

# Cisco Compute Security Overview

Version 2.0

March 2025

---

# Contents

Intended use and audience	4
Bias statement	4
Legal notices	4
Prerequisites	4
Introduction	4
Six pillars of security	6
1. The Cisco Secure Development Lifecycle–CSDL	7
<b>CSDL philosophy</b>	7
<b>Development milestones</b>	8
<b>CSDL product adherence methodologies</b>	8
<b>Cisco Security and Trust Organization</b>	9
Advisories, vulnerabilities, and incident responses	10
<b>CERT advisories</b>	10
<b>Additional vulnerability testing measures</b>	10
<b>Incident response</b>	10
2. Supply chain security	11
<b>Counterfeit prevention</b>	11
<b>Consortiums for secure vendors</b>	12
3. Certifications and compliance	13
<b>Certification process</b>	13
<b>Common criteria</b>	13
<b>SOC 2 Type 2</b>	13
<b>FIPS</b>	15
<b>FIPS 140-3</b>	16
<b>IPv6</b>	17
<b>Defense Information Security Agency Approved Product List</b>	17
<b>Other certifications and procedural guidelines</b>	17

---

4. The mechanics of server security – system-level security	18
<b>Physical security</b>	<b>18</b>
<b>Card boot - TAM</b>	<b>18</b>
<b>System boot</b>	<b>19</b>
<b>Deployment and management at scale</b>	<b>26</b>
<b>SSL key management: UI certificates and self-encrypting drives</b>	<b>34</b>
5. Secure application operation	42
<b>Confidential computing</b>	<b>42</b>
6. Secure data delivery and storage	45
<b>Encryption and key management</b>	<b>45</b>
<b>Self-encrypting drives</b>	<b>46</b>
<b>Virtual interface card</b>	<b>48</b>
Conclusion	48
For more information	49
Document information	50

---

## Intended use and audience

This document contains confidential material that is proprietary to Cisco Systems, Inc. The materials, ideas, and concepts contained herein are to be used exclusively to assist in the configuration of Cisco corporation's hardware and software solutions.

## Bias statement

The documentation set for this product strives to use bias-free language. For purposes of this documentation set, bias-free is defined as language that does not imply discrimination based on age, disability, gender, racial identity, ethnic identity, sexual orientation, socioeconomic status, and intersectionality. Exceptions may be present in the documentation due to language that is hardcoded in the user interfaces of the product software, language used based on standards documentation, or language that is used by a referenced third-party product.

## Legal notices

All information in this document is provided in confidence and shall not be published or disclosed, wholly or in part to any other party without Cisco's written permission.


## Prerequisites

We recommend reviewing the Cisco UCS® release notes, installation guide, and user guide before proceeding with any configuration. Please contact Cisco® Technical Assistance Center (Cisco TAC) or your Cisco representative if you need assistance.

## Introduction

In the continuously evolving server landscape, Cisco Systems, Inc., has consistently developed solutions to address critical user concerns regarding security, performance, and reliability. This white paper serves as a holistic solution overview detailing the robust security posture within the Cisco Unified Computing System™ (Cisco UCS) platform. This guide explores the development philosophy, industry-leading certifications, robust management and scalable deployment features, confidential computing mechanisms, and secure storage capabilities present within the Cisco UCS ecosystem.

At the core of Cisco's UCS platform lies a development philosophy centered on proactive security measures, with an approach designed for preemptive threat mitigation and continuous enhancement.



**Figure 1.**  
The Cisco security value chain

Cisco leverages in-house technologies and research to fortify its UCS architecture against emerging threats. Incorporating robust industry practices and adhering to stringent security protocols, the UCS platform is built to meet the highest standards of security certifications, ensuring compliance with regulatory frameworks, and assuring customers of a resilient and safeguarded infrastructure. Moreover, the management features embedded within the UCS solution provide administrators with comprehensive tools for monitoring, auditing, and controlling access, enabling proactive threat identification and rapid response to potential security breaches.

Throughout the building process, Cisco UCS has the distinct advantage of full stack development. This encompasses a secure supply chain and vetted vendors and extends through the manufacturing process through programs such as anti-counterfeiting, to the entire development lifecycle. The benefits of this full-stack process are mandated secure practices at every step of the build and development process for Cisco UCS.

In addition to its development and certification framework, Cisco UCS utilizes advancements in confidential computing and secure storage to keep user applications and data protected. Implementing NIST-approved encryption techniques, secure boot processes, and hardware-based isolation mechanisms, UCS ensures data confidentiality, integrity, and availability throughout its lifecycle. Through secure storage solutions, and federally certified secure interfaces, users leverage the UCS platform confidently, knowing that their data remains protected against unauthorized access. This white paper discusses the implementation of these features, demonstrating how UCS meets and exceeds the security and accountability requirements in today’s enterprise environments.

This approach allows Cisco UCS to extend its secure reach into the deployment process. This is accomplished by the strict and segmented enforcement of security and configuration policies that are built into Cisco UCS Manager (UCSM) and Cisco Intersight®. These capabilities ensure that there is no configuration violation or drift, resulting in a robust and transparent ability to verify compliance with security processes and to monitor the outcome.

---

All of these capabilities serve to guarantee that the Cisco UCS system and its deployment adhere to a zero-trust security framework. A zero-trust security framework is a comprehensive approach to security that assumes no user, system, or device can be trusted by default, regardless of its location relative to the network perimeter, or whether the compute system has previously verified the user. It operates under the principle of "never trust, always verify," meaning that every access request is thoroughly verified before granting access, irrespective of where it originates. The zero-trust framework strives to protect modern digital environments by leveraging policies, network segmentation, preventing lateral movement (for example, configuration drift), providing Layer-7 threat prevention and simplifying granular user-access control.

Implementing a zero-trust framework provides following additional benefits:

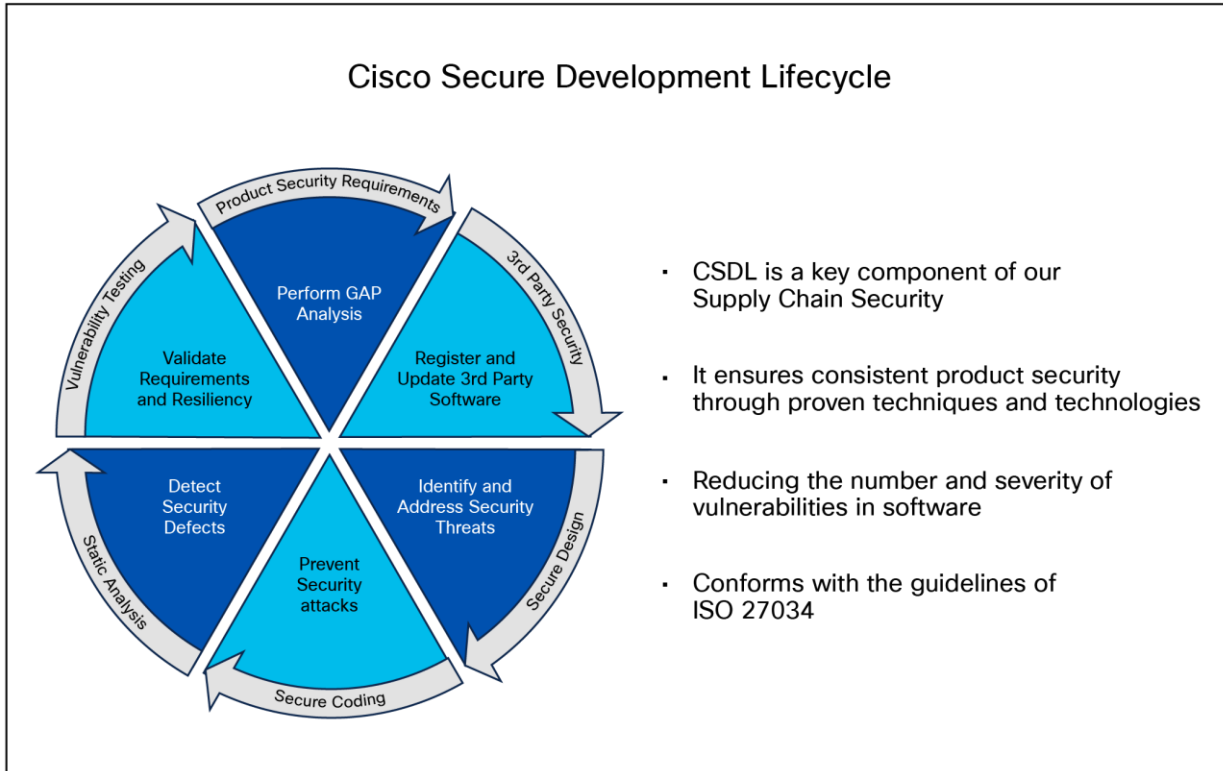
- **Enhanced security:** By treating every access request or deployment change as a potential threat, zero trust significantly reduces the risk of data breaches and other security incidents.
- **Greater visibility:** Constant monitoring of system activities provides a comprehensive view of the network, enabling quick identification and response to any unusual or suspicious activities.
- **Reduced attack surface:** By enforcing least privilege access and micro-segmentation, zero trust minimizes the potential points of vulnerability in the system.
- **Improved compliance:** The stringent security controls in zero trust can help organizations meet compliance requirements for data protection and privacy.
- **Efficient incident response:** Quickly detect, block, and respond to threats, verify data integrity, and implement data loss prevention.
- **Protection against internal threats:** Zero trust considers the possibility of threats coming from inside the organization or network, offering protection against insider threats as well as external ones.

## Six pillars of security

Security in the compute realm can be conveniently divided into six subcategories. First, we will look at the secure development process for the UCS platform; second, we will examine secure supply chains and manufacturing, because these two are the foundation for a secure product. Next, we dive into the certification landscape for the product. Then we investigate the mechanics of server security at the system level. This will encompass secure boot, deployment, management, and monitoring. Then we will examine the features available for the secure operation of applications. This will involve practices around confidential computing and hardware-based isolation and encryption. Finally, we will examine secure data storage and delivery. This will focus on self-encrypting drives and network delivery.

# 1. The Cisco Secure Development Lifecycle–CSDL

Cisco UCS products and components are developed, integrated, and tested using the Cisco Secure Development Lifecycle (CSDL). Secure product development and deployment has several components, ranging from following specified design and development practices, to testing their implementation, to providing customers with a set of recommendations for deployments that maximize the security of the system.



**Figure 2.**  
The Cisco Secure Product Development Lifecycle

## CSDL philosophy

A poor product design can open the way to vulnerabilities. The CSDL is designed to mitigate these potential issues. At Cisco, our secure-design approach requires two types of considerations:

- Design with security in mind
- Use threat modeling to validate the design's security

Designing with security in mind is an ongoing commitment to personal and professional improvement through:

- Training
- Applying Product Security Baseline (PSB) design principles
- Considering other industry-standard secure-design principles
- Being aware of common attack methods and designing safeguards against them
- Taking full advantage of designs and libraries that are known to be highly secure
- Protecting all potential entry points

---

Cisco also reduces design-based vulnerabilities by considering known threats and attacks:

- Follow the flow of data through the system
- Identify trust boundaries where data may be compromised
- Based on the data flow diagram, generate a list of threats and mitigations from a database of known threats, tailored by product type
- Prioritize and implement mitigations to the identified threats

The goal of this effort is to enforce a set of security processes and ensure a security mind set at every stage of development:

- Secure design
- Secure coding
- Secure analysis
- Vulnerability testing
- Secure deployments

Cisco UCS product development focuses on two areas to satisfy the CSDL model:

- Internal requirements
  - Adhere to the secure-development process
- Market-based requirements
  - Complete and validate against certifications (federal)
  - Document and educate

## Development milestones

Each iteration of the product's development addresses needs for ongoing security fixes and general feature enhancements that include security components (new deployment models, changes in management, partner onboarding, etc.). At every stage of development, the product(s) undergo potential enhancements relative to findings and new features.

- The system is configured in QA to accommodate the relevant settings identified above and run through a typical deployment test.
- The result is a validated set of best practices for security and is communicated through the CSDL process and exposed in the documentation.

## CSDL product adherence methodologies

Cisco CSDL adheres to Cisco Product Development Methodology (PDM), ISO/IEC 27034, and ISO 9000 compliance requirements. The ISO/IEC 27034 standard provides an internationally recognized standard for application security. Details for ISO/IEC 27034 can be found [here](#). The ISO 9000 family of quality management systems standards is designed to help organizations ensure that they meet the needs of customers and other stakeholders while meeting statutory and regulatory requirements related to a product or service. ISO 9000 details are [here](#).

The CSDL process is not a one-time approach to product development. It is recursive, with vulnerability testing, penetration testing, and threat modelling added to subsequent development of CSDL. This process follows ISO

9000 and ISO 27034 standards as part of an internationally recognized set of guidelines. The approaches involved often use a solution-wide methodology; for example, utilizing our continually updated CiscoSSL crypto module to guarantee that Cisco UCS (along with other elements in the Cisco offering) is always secure and meets FIPS certification requirements.

### Cisco Security and Trust Organization

Cisco Security and Trust Organization (S&TO) has the core responsibility to implement CSDL. In the effort to accomplish this, S&TO encompasses various groups with core responsibilities to deliver a secure product or respond to security concerns as they arise.



**Figure 3.**  
The groups within Cisco S&TO

---

## Advisories, vulnerabilities, and incident responses

### **CERT advisories**

Cisco's Computer Emergency Response Team (CERT) advisories are transmitted when new vulnerabilities are identified. Cisco's internal CERT team monitors and alerts product groups to potential issues that might affect their respective components. When these items are identified by CERT or are otherwise indicated by vendor partners (VMware, etc.), patches are either developed or acquired from the respective vendors. Cisco has heavily invested to protect customers by creating this team, which constantly monitors threats and builds a centralized solution to remediate these issues and vulnerabilities.

### **Additional vulnerability testing measures**

Cisco also utilizes an internal tool for threat modeling called Threat Builder. This tool is used to explicitly map out application components and services and to identify potential attack surfaces and develop line items for direct evaluation. This information, along with industry tools, is used for vulnerability and exploit testing by Cisco's ASIG (Advanced Security Initiatives Group). ASIG also uses fuzzing and manual testing as part of its suite of tools.

### **Incident response**

The Cisco Product Security Incident Response Team (PSIRT) is responsible for responding to Cisco product security incidents. The Cisco PSIRT is a dedicated, global team that manages the receipt, investigation, and public reporting of information about security vulnerabilities and issues related to Cisco products and services. Cisco defines a security vulnerability as a weakness in the computational logic (for example, code) found in software and hardware components that, when exploited, results in a negative impact to confidentiality, integrity, or availability. Cisco reserves the right to deviate from this definition based on specific circumstances. The Cisco PSIRT adheres to ISO/IEC 29147:2018, which is a set of [guidelines for disclosure of potential vulnerabilities](#) established by the International Organization for Standardization.

The Cisco PSIRT is on call and works 24 hours a day with Cisco customers, independent security researchers, consultants, industry organizations, and other vendors to identify possible security vulnerabilities and issues with Cisco products and networks.

All vulnerabilities disclosed in Cisco Security Advisories are assigned a Common Vulnerability and Exposures (CVE) identifier and a Common Vulnerability Scoring System (CVSS score) to aid in identification. Additionally, all vulnerabilities are classified based on a Security Impact Rating (SIR).

Cisco uses version 3.1 of CVSS as part of its standard process of evaluating reported potential vulnerabilities in Cisco products. The CVSS model uses three distinct measurements or scores that include base, temporal, and environmental calculations. Cisco provides an evaluation of the base vulnerability score and, in some instances, a temporal vulnerability score. End users are encouraged to compute the environmental score based on their network parameters.

---

In addition, Cisco uses the Security Impact Rating (SIR) as a way to categorize vulnerability severity in a simpler manner. The SIR is based on the CVSS base score, adjusted by PSIRT to account for variables specific to Cisco, and is included in every Cisco Security Advisory.

Cisco PSIRT assigns a Common Vulnerabilities and Exposures Identifier (CVE ID) to any vulnerability that is found in Cisco products and that qualifies to receive this identifier. Usually, all vulnerabilities with medium, high, or severe SIRs—that is, a CVSS score of 4.0 or greater—will qualify for a CVE ID.

## 2. Supply chain security

A critical aspect of secure product development and deployment is ensuring that the components that go into the system are legitimate and uncompromised. To this end, Cisco takes exceptional measures to ensure supply chain integrity.

### Counterfeit prevention

The Cisco Value Chain Security Program (Value Chain describes the development model used for all Cisco products. Cisco is a leader in industry and international standards on counterfeit reduction and has been engaged in decades-long efforts to prevent and detect the distribution of counterfeit products. Cisco incorporates tools and processes to prevent counterfeiting—beginning with product development, through the manufacturing process, and in the marketplace. The anti-counterfeiting methodology implemented into Cisco UCS products guarantees that customers are getting authentic hardware that has not been manipulated or altered. The system has secure roots of trust built into the hardware to specifically perform certificate checks to verify that the hardware was manufactured by Cisco.

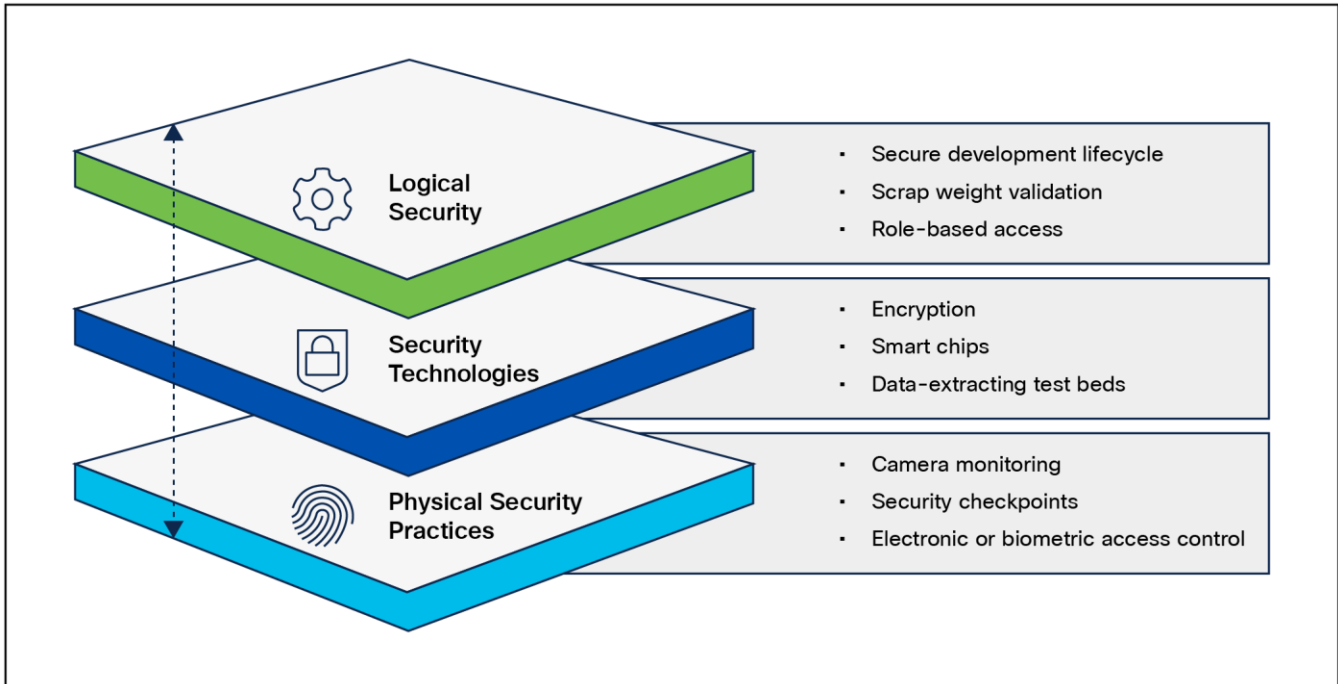
In collaboration with Cisco's Brand Protection and legal teams, and other Cisco teams, an end-user portal has been developed to aid customers in these efforts and can be accessed at: [anticounterfeit.cisco.com](https://anticounterfeit.cisco.com).

Cisco's Brand Protection has conducted numerous investigations into counterfeiting operations and worked with local law enforcement to disrupt those operations. The portal includes examples of the Brand Protection Team's work over the years, and the numerous resources that are available for Cisco customers and partners.

This Value Chain has the following characteristics:

- Comprehensive across all stages of solution's lifecycle
- Multilayer approach, with protection focused against:
  - Source-code corruption
  - Hardware counterfeit
  - Misuse of intellectual property

This multilayered approach is shown below.



**Figure 4.**  
Layers of the Cisco Value Chain

### Consortiums for secure vendors

The table below shows a list of the secure supply chain consortiums of which Cisco is a member. These consortiums help guide industry standards on securing component access, acquisition, manufacturing, and delivery. Membership in these groups gives Cisco a voice in developing standards and setting precedents in this secure space.

**Table 1.** Secure supply chain consortiums

Name	Component(s)	Description	Status
<b>Transported Asset Protection Association (TAPA)</b>	Supply chain	The Transported Asset Protection Association's (TAPA) security standards act as a worldwide benchmark for supply chain security and resilience, providing guidance, processes, and tools that reduce loss exposure, protect assets, and reduce the costs of cargo theft.	Member
<b>Customs Trade Partnership Against Terrorism (CTPAT)</b>	Supply chain	Customs Trade Partnership Against Terrorism (CTPAT) Trade Compliance program is a voluntary program that provides the opportunity for importers who have made a commitment of resources to assume responsibility for monitoring their own compliance in exchange for benefits.	Member

---

## 3. Certifications and compliance

### Certification process

Federal compliance and audit-based certifications are a critical component of a standardized and predictable security posture. They are critical in most federal deployments, especially those dealing with financial and defense arenas. The Cisco Global Certification Team (GCT) works to complete various certifications.

### Common criteria

Common Criteria for Information Technology Security Evaluation (Common Criteria or CC) is an International Standard (ISO/IEC 15408) for computer security certification, currently in v3.1 rev 5.

- System users specify their security functional and assurance requirements through the use of protection profiles; vendors can then make claims about the security attributes of their products, and testing laboratories can evaluate the products to determine if they actually meet the claims, following a process like the following:
  - My customers have some security needs defined in a set of CC guidelines.
  - This is my system; this is what I say it can do to meet those security needs.
  - Let us (vendor and lab) agree on a test; here is the procedure.
  - Here are my results.
  - You (lab) run it on your own and verify those results.
  - If successful, deliver certification.

A key part of Evaluation Assurance Level (EAL) is the Security Target document, which comprises a rigorous definition of functions, features, and intended use, tailored for the specific hardware or software component under test (the TOE, Target of Evaluation). The EAL rating determines the extent of the testing, and the confidence that security is as claimed. Simply, EAL indicates the degree to which something does what it says it does.

### SOC 2 Type 2

System and Organization Controls (SOC) (also sometimes referred to as service organizations controls), as defined by the [American Institute of Certified Public Accountants](#) (AICPA), is the name of a suite of reports produced during an audit. SOC is intended for use by service organizations (organizations that provide information systems as a service to other organizations) to issue validated reports on [internal controls](#) over those information systems to the users of those services.

Generally, SOC 1, SOC 2, or SOC 3 are referred to regarding compliance; however, SOC for cybersecurity and SOC for supply chain certifications also exist. SOC compliance and audits are intended for organizations that provide services to other organizations. For example, a company that offers cloud-hosting services may need SOC compliance. For Cisco and its customers, this is relevant for the Cisco Intersight SaaS cloud service.

There are two levels of SOC reports that are also specified by SSAE 18:

- Type 1, which describes a service organization's systems and whether the design of specified controls meet the relevant trust principles.
- Type 2, which also addresses the operational effectiveness of the specified controls over a period of time (usually 9 to 12 months).

---

There are three types of SOC reports:

- SOC 1 – Internal Control over Financial Reporting (ICFR)[4]
- SOC 2 – Trust Services Criteria[5][6]
- SOC 3 – Trust Services Criteria for General Use Report[7]

SOC 2 Type 2 certified: meets controls for confidentiality, security, and availability, among others.

SOC 2 reports focus on controls addressed by five semi-overlapping categories called Trust Service Criteria, which also support the CIA triad of information security:

1. Security – information and systems are protected against unauthorized access and disclosure, and damage to the system that could compromise the availability, confidentiality, integrity and privacy of the system.
  - a. Firewalls
  - b. Intrusion detection
  - c. Multifactor authentication
2. Availability – information and systems are available for operational use.
  - a. Performance monitoring
  - b. Disaster recovery
  - c. Incident handling
3. Confidentiality – information is protected and available on a legitimate need-to-know basis. Applies to various types of sensitive information.
  - a. Encryption
  - b. Access controls
  - c. Firewalls
4. Processing integrity – system processing is complete, valid, accurate, timely, and authorized.
  - a. Quality assurance
  - b. Process monitoring
  - c. Adherence to principle
5. Privacy – personal information is collected, used, retained, disclosed, and disposed according to policy. Privacy applies only to personal information.
  - a. Access control
  - b. Multifactor authentication
  - c. Encryption

---

## FIPS

The Federal Information Processing Standard (FIPS) Publication 140-2 is a U.S. government computer-security standard used to approve cryptographic modules.

Cisco UCS is compliant with FIPS140-2 level 1 through direct implementation of the FIPS-compliant CiscoSSL crypto module. The module, once implemented, is vetted by a third party that is federally certified to ascertain compliance status.

- Utilizes CiscoSSL module
  - Already FIPS compliant
  - SSH-approved cipher list
  - SSL/TLS implementation
  - Eliminates weak or compromised components
- Regularly updated
- Lab validates that the module is incorporated correctly.
  - Build logs
  - Source-access-identifying calls to the module
  - All admin access points to the cluster are covered here.
- SSH for CLI
- HTTPS for UI

A comprehensive list of Cisco FIPS-compliant products is given below, along with the corresponding reference with NIST:

- Cisco FIPS-certified products: <http://www.cisco.com/c/en/us/solutions/industries/government/global-government-certifications/fips-140.html>.
- Cryptographic Module Validation Program (CMVP) vendor list: [Cryptographic Module Validation Program | CSRC \(nist.gov\)](#).

An official website of the United States government. [Here's how you know](#)

**NIST** Information Technology Laboratory  
**COMPUTER SECURITY RESOURCE CENTER**

PROJECTS CRYPTOGRAPHIC MODULE VALIDATION PROGRAM VALIDATED MODULES

## Cryptographic Module Validation Program CMVP

f t in e

### Search

All questions regarding the implementation and/or use of any validated cryptographic module should first be directed to the appropriate VENDOR point of contact (listed for each entry). General CMVP questions should be directed to [cmvp@nist.gov](mailto:cmvp@nist.gov).

Use this form to search for information on validated cryptographic modules.

Select the basic search type to search modules on the active validation list. Select the advanced search type to search modules on the historical and revoked module lists.

Search Type:  Basic  Advanced

Certificate Number:

Vendor:

**Figure 5.**  
FIPS vendor listings

## FIPS 140-3

FIPS 140-3 is the successor to FIPS 140-2. FIPS 140-3 became effective on September 22, 2019. FIPS 140-3 testing began on September 22, 2020, and a small number of validation certificates have been issued. FIPS 140-2 testing was available until September 21, 2021, creating an overlapping transition period of one year.

FIPS 140-2 test reports that remain in the CMVP queue will still be granted validations after that date, but all FIPS 140-2 validations will be moved to the Historical List on September 21, 2026, regardless of their actual final validation date. Versions of Cisco software using CiscoSSL version 8.3 or later are certified for FIPS 140-3 encryption.

Note that CiscoSSH uses the cryptographic engine from CiscoSSL so that it is automatically covered. This includes the latest UCSM and CIMC software that manages the Device Connector for communication to Intersight SaaS as well as Intersight appliances such as the PVA (Private Virtual Appliance) and the CVA (Connected Virtual Appliance). The certification letter can be found here ([Cryptographic Module Validation Program | CSRC](#)).

CNSA (Commercial National Security Algorithm) is a schema that is detailed in RFC 9151: [RFC 9151: Commercial National Security Algorithm \(CNSA\) Suite Profile for TLS and DTLS 1.2 and 1.3 \(rfc-editor.org\)](#).

It describes which algorithms should be in use and what their profiles should look like. It is intended to give guidance for secure and interoperable communications, including guidelines for certificates, for national security reasons.

---

Cisco supports both Elliptic Cryptographic Certificates (ECC) and RSA certificates, so this requirement is met:

- “Elliptic Curve Digital Signature Algorithm (ECDSA) and Elliptic Curve Diffie–Hellman (ECDH) key pairs are on the curve P-384. FIPS 186-4, Appendix B.4, provides useful guidance for elliptic curve key pair generation that SHOULD be followed by systems that conform to the RFC.
- RSA key pairs (public, private) are identified by the modulus size expressed in bits; RSA-3072 and RSA-4096 are computed using moduli of 3072 bits and 4096 bits, respectively Cisco’s FIPS certification through CiscoSSL implements federally approved crypto modules to satisfy the complexity requirements as well.

CNSA compliance is just a matter of making sure to implement a cryptographic ecosystem according to the CNSA requirements since Cisco UCS supports all the documented methods.

## IPv6

The U.S. Office of Management and Budget (OMB) has directed [OMB-2020, OMB-2010, and OMB-2005] the National Institute of Standards and Technology (NIST) to develop the technical infrastructure (standards and testing) necessary to support wide-scale adoption of IPv6 in the U.S. Government (USG). In response, NIST developed a technical standards profile for U.S. government acquisition of IPv6-enabled networked information technology. The USGv6 profile includes a forward-looking set of protocol specifications published by the Internet Engineering Task Force (IETF), encompassing basic IPv6 functionality, and specific requirements and key optional capabilities for routing, security, multicast, network management, and quality of service.

The profile also contains NIST-defined requirements for IPv6-aware firewalls and intrusion-detection systems. The program also established a robust testing infrastructure to enable IPv6 products to be tested for compliance to profile requirements and for interoperability by accredited laboratories using standardized test methods. Cisco UCS platforms are in the process of completing this qualification.

## Defense information security agency approved product list

The Defense Information Security Agency Approved Product List is a multifaceted U.S. federal certification that gives approval for products to operate in secure environments. Certification is currently being sought for Cisco with the Cisco Global Certification Team and the Cisco Compute Business Unit.

## Other certifications and procedural guidelines

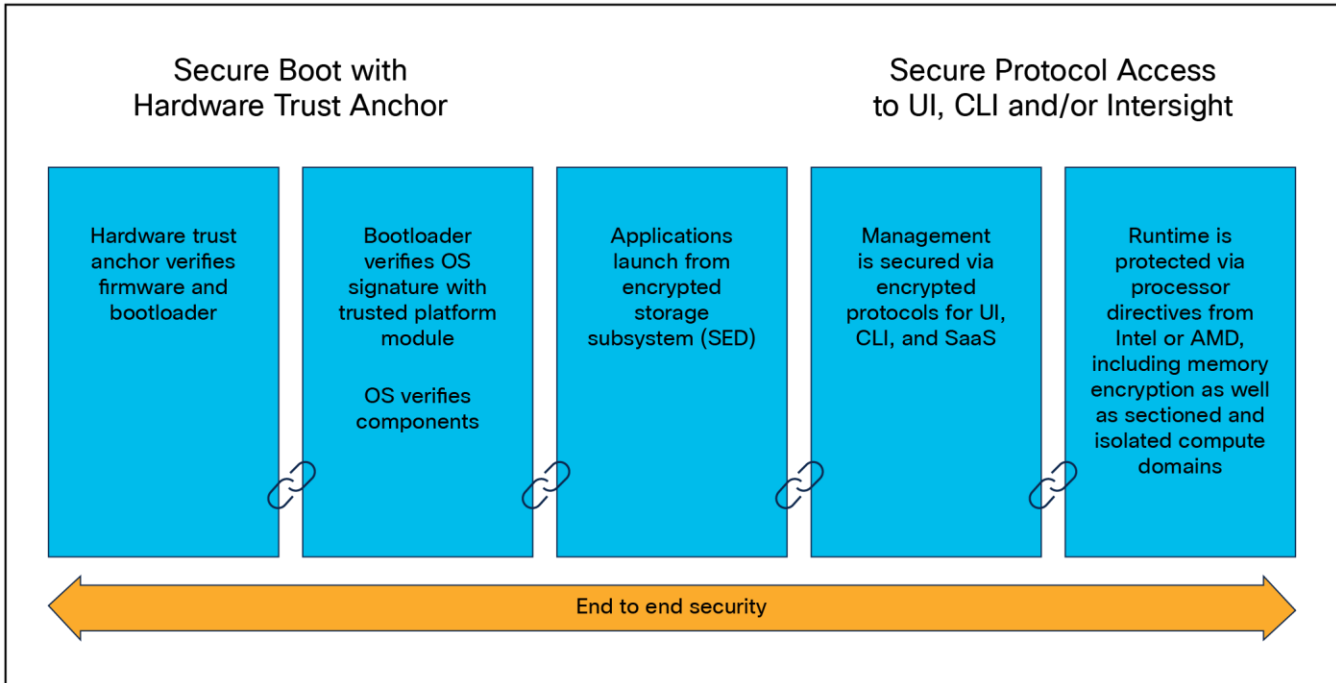
ISO/IEC 27001 is not a certification for specific pieces of hardware as much as it is a dozen or so “best practices” in the form of checklists and guidelines for how organizations manage their security controls internally. It observes such things as building access, password management, badging into a copier to make copies, etc. Training on a frequent basis is a part of the standard.

Cisco is ISO/IEC 27001 certified. This is a link to our ISO/IEC 27001 certificate [Cisco Secure Cloud Analytics \(StealthWatch\) ISO/IEC 27001:2013, 27017:2015, 27018:2019](#).

ISO/IEC 27001:2013: the Cisco Intersight platform has completed its ISO/IEC 27001:2013 First Surveillance Audit from the external certification body/auditor Coalfire, and the certificate issued has been uploaded to [Trust Portal site](#). The First Surveillance Audit included a review of the establishment and overall operating effectiveness of control areas that form Cisco Intersight’s Information Security Management System.

## 4. The mechanics of server security – system-level security

The first pillar we will examine is the mechanics of server security at the system level. This encompasses ensuring that the system is running legitimate software, is free from counterfeit hardware, and is a generally trustworthy platform for handling user data. This is part of Cisco’s end-to-end security philosophy with respect to server system level security.



**Figure 6.** End-to-end security at the system level from bootloader to runtime defenses

### Physical security

Physical security is an important aspect of general security for Cisco UCS systems. As such, Cisco UCS servers incorporate chassis intrusion detection as part of their security features. While the exact implementation details may vary, the concept is straightforward. A microswitch is typically placed on the server chassis so that when the chassis is opened, the switch is released. The switch is connected to the motherboard and power [is supplied by the CMOS battery, allowing the detection system to work even when the server is unplugged.](#)

### Card boot - TAM

Cisco UCS systems have a variety of add-in cards that serve many different functions. These range from additional VICs, to NICs, offload DPUs, GPUs, and various HBAs. As part of a secure deployment posture, it is not only important to be able to securely boot the main system, including UEFI (Unified Extensible Firmware Interface) BIOS, bootloader, and operating system. The system must also securely boot the firmware that runs on the add-in cards themselves. To this end, most cards in use with Cisco UCS have a trust anchor built in that validates the card BIOS at card boot. These systems serve the same function as the TAM on the Cisco UCS motherboard in securing server firmware and ensuring legitimate code execution in system start-up.

## System boot

A secure system boot relies on a set of trusted Cisco technologies. Here are the fundamental concepts of Cisco Trustworthy Technologies:

### Chain of trust

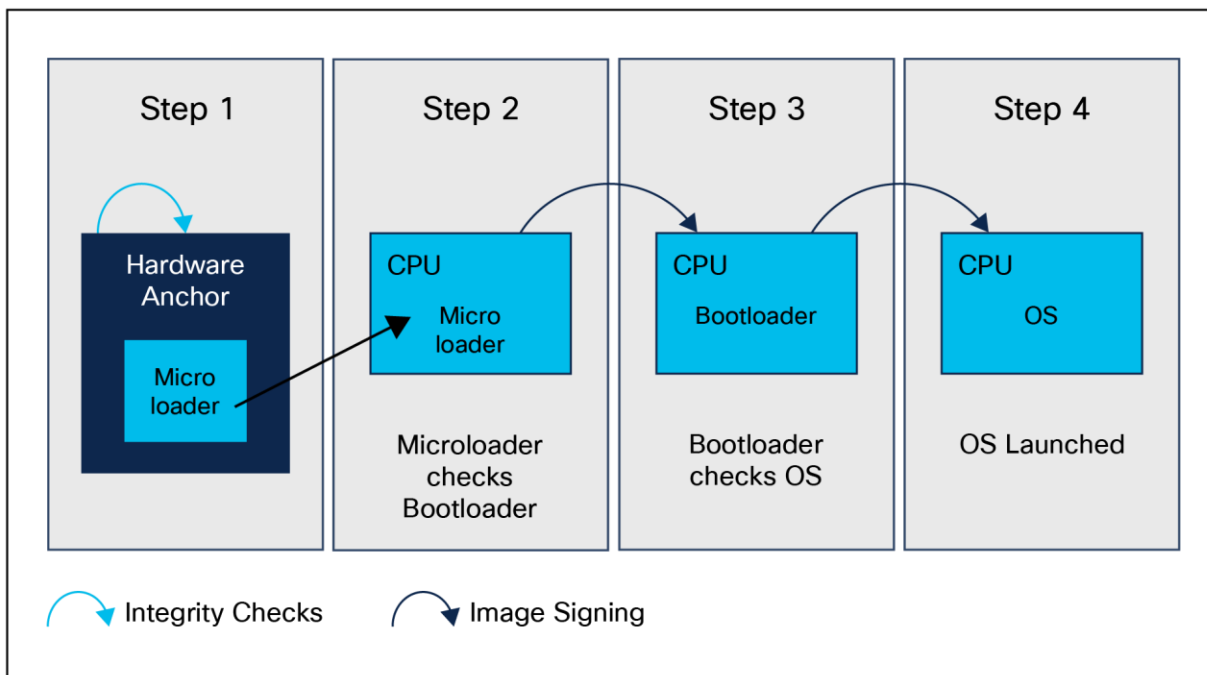
A chain of trust exists when the integrity of each element of code in a system is validated before that piece of code is allowed to run. A chain of trust starts with a root of trust. The root of trust validates the next element in the chain (usually firmware) before it is allowed to start, and so on. Through the use of image signing and trusted elements, a chain of trust can be created that boots the system securely and validates the integrity of Cisco software.

### Cisco Secure Boot

Cisco Secure Boot helps to ensure the code that executes on Cisco hardware platforms is authentic and unmodified. A Cisco hardware-anchored secure boot protects the microloader (the first piece of code that boots) in tamper-resistant hardware, establishing a root of trust that helps prevent Cisco compute and network devices from executing illegitimate or malicious software. A subsequent boot of the installed operating system is verified and attested with the Trusted Platform Module (TPM).

Cisco Secure Boot helps ensure that the code that executes on Cisco hardware platforms is genuine and untampered. A typical UEFI-based boot process starts at the UEFI firmware and works up to the bootloader and the operating system. A tampered UEFI firmware can result in the entire boot process being compromised.

Using a hardware-anchored root of trust, digitally signed software images, and a unique device identity, Cisco hardware-anchored secure boot establishes a chain of trust which boots the system securely and validates the integrity of the software. The root of trust (a.k.a. the microloader), which is protected by tamper-resistant hardware, first performs a self-check and then verifies the UEFI firmware, thus beginning a chain of trust leading up to the integrity verification of the entire Cisco UCS operating system.



**Figure 7.**  
Cisco Secure Boot process

---

## Image signing

Image signing is a two-step process for creating a unique digital signature for a given block of code. First, a hashing algorithm, similar to a checksum, is used to compute a hash value of the block of code. The hash is then encrypted with a Cisco private key, resulting in a digital signature that is attached to and delivered with the image. Signed images may be checked at runtime to verify that the software has not been modified.

## Hardware root of trust – the Trust Anchor module and the Trusted Platform module (2.0).

A trusted element in system software is a piece of code that is known to be authentic. A trusted element must either be immutable (stored in such a way as to prevent modification) or authenticated through validation mechanisms. Cisco anchors the root of trust, which initiates the boot process, in tamper-resistant hardware. The hardware-anchored root of trust protects the first code running on a system from compromise and becomes the root of trust for the system.

The Trust Anchor module (TAM) is a proprietary, tamper-resistant chip found in many Cisco products and features nonvolatile secure storage, a Secure Unique Device Identifier (SUDI), and crypto services such as Random Number Generation (RNG), secure storage, key management, and crypto services to the running OS and applications.

The hardware root of trust is a Cisco ACT2 Trust Anchor module (TAM). This module has the following characteristics:

- Immutable identity with IEEE 802.1AR (Secure UDI- X.509 cert)
- Anti-theft and anti-counterfeiting
- Built-in cryptographic functions
- Secure storage for certificates and objects
- Certifiable NIST SP800-92 random number generation

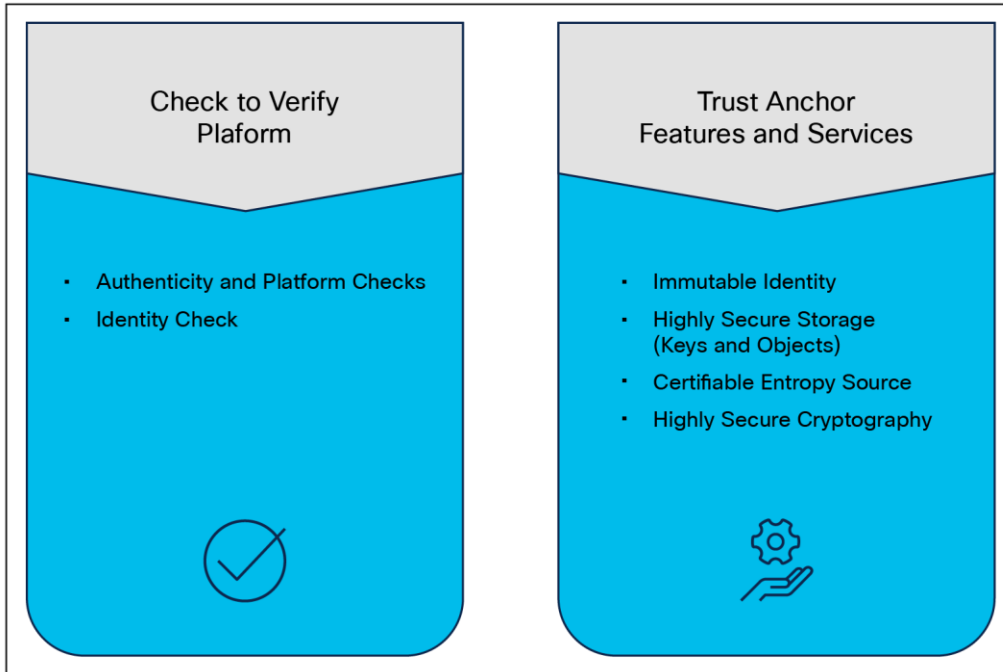
Once a system is securely booted, it is often important to get external verification that this is indeed the case. This is done through attestation. “Attestation” is evidence of a result; for example, “The host was booted with secure boot enabled and signed code.” This is accomplished through the Trusted Platform module (TPM).

## Immutable identity

The Secure Unique Device Identifier (SUDI) is an X.509v3 certificate that maintains the product identifier and serial number. The identity is implemented at manufacture and is chained to a publicly identifiable root certificate authority. The SUDI can be used as an unchangeable identity for configuration, security, auditing, and management.

The SUDI credential in the Trust Anchor module can be based on either RSA or Elliptic Curve Digital Signature Algorithm (ECDSA). The SUDI certificate, the associated key pair, and its entire certificate chain are stored in a tamper-resistant Trust Anchor module chip. Furthermore, the key pair is cryptographically bound to a specific Trust Anchor chip, and the private key is never exported. This feature makes cloning or spoofing the identity information virtually impossible.

The SUDI can be used for asymmetric key operations such as encryption, decryption, signing, and verifying that allow passage of the data to be operated upon. This capability makes remote authentication of a device possible. It enables accurate and consistent electronic identification of Cisco products for asset management, provisioning, version visibility, service entitlement, quality feedback, and inventory management.



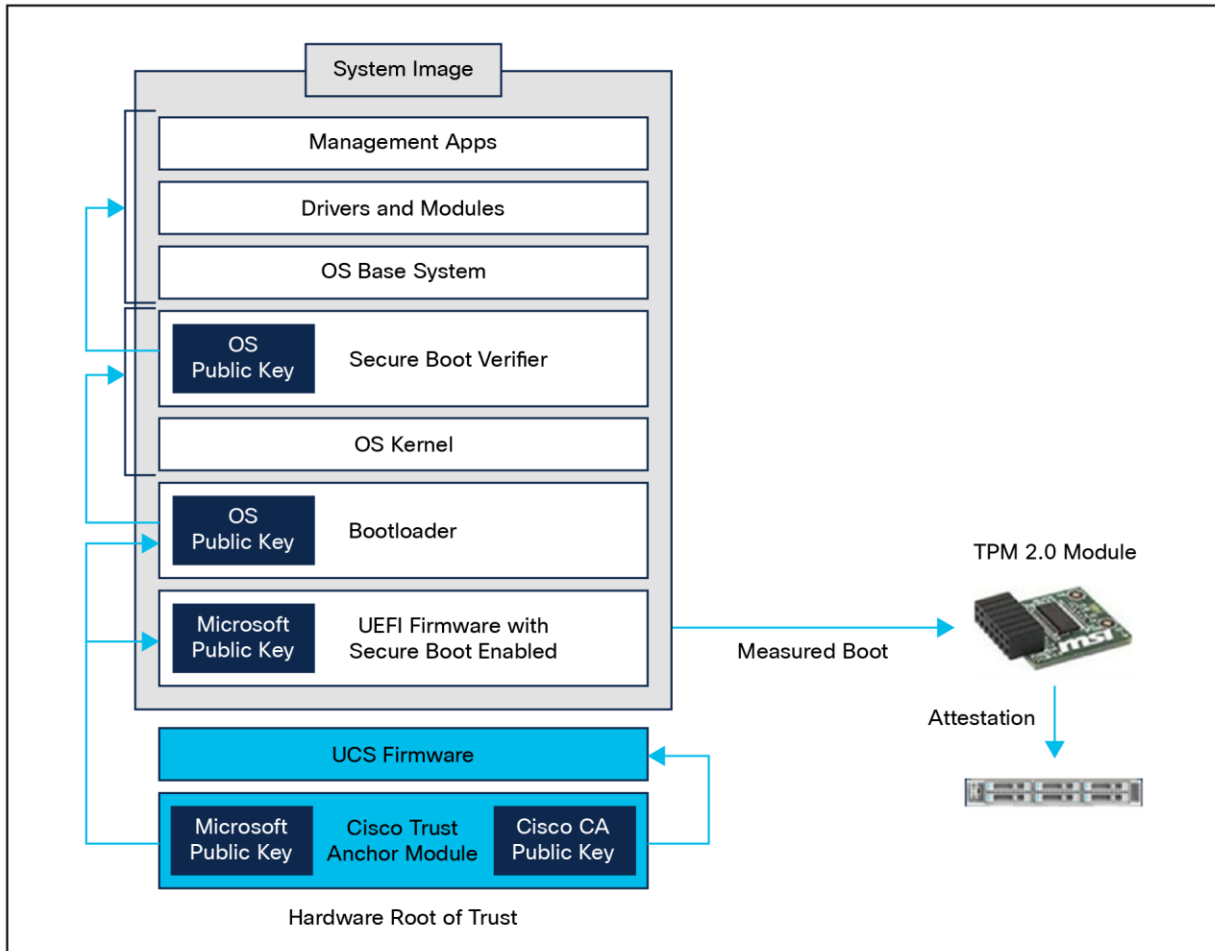
**Figure 8.**  
TAM functions

Currently, the secure boot process, when enabled, is in effect during boot of both the system firmware and the installed operating system. The end-to-end security model that this enables, when combined with the secure UI and CLI, includes the hardware Trust Anchor module (TAM), which is used to secure the system boot, in order to secure OS boot with externally verifiable attestation using the Trusted Platform module (TPM).

This implementation covers the following:

- Uses secure boot, which is secured by public keys stored in the write-protected hardware root of trust
- Ensures that only a trusted OS image, including drivers, is booted by verifying signatures
- Supports attestation of secure boot through TPM 2.0

The detailed process flow for secure boot of the system and OS with attestation capability is shown below. Note that the certificate-based hardware root of trust validates the Cisco UCS firmware, which ensures that a clean BIOS is set for key validation of the hypervisor bootloader, and so on. This guarantees that the hardware and operating system in the UCS system has not been tampered with. External validation of this can be made through attestation using the TPM 2.0 module in Cisco UCS.



**Figure 9.**  
Use of TAM and TPM in the entire process

### Runtime defenses

Runtime Defenses (RTDs) target injection attacks of malicious code in running software. Cisco runtime defenses include Address Space Layout Randomization (ASLR), Built-in Object Size Checking (BOSC), and X-space.

RTDs have the following features:

- Runtime defenses make it harder or impossible for attackers to exploit vulnerabilities in running software.
- They are complementary; you can implement them individually or deploy several runtime defenses together.

### CPU hardware protections

Cisco UCS supports both Intel® and AMD processors. The latest generations of these CPUs and their accompanying chipsets have extensions and programmatic capabilities regarding memory encryption and secure code execution and isolation.

---

## Intel Boot Guard

Intel Boot Guard (4th gen CPU and greater) is a security technology designed to enhance the integrity of the boot process and protect against unauthorized firmware and bootloader modifications on systems that use Intel processors. It is part of Intel's broader security initiatives to safeguard the boot process from potential threats and to ensure that the system starts up securely. Here are the key aspects of Intel Boot Guard:

- **Boot process integrity:**
  - Intel Boot Guard focuses on protecting the boot process, ensuring that the system starts up using only authorized and unaltered firmware and bootloader components.
- **Hardware-based protection:**
  - Boot Guard operates at the hardware level, utilizing a combination of hardware-based mechanisms within the Intel chipset and processor.
- **Verified boot:**
  - During the boot process, Boot Guard verifies the digital signature of the firmware and bootloader components before allowing them to execute. Digital signatures are used to verify the authenticity and integrity of the firmware and bootloader code.
- **Measures against unauthorized modifications:**
  - Boot Guard helps prevent unauthorized modifications to the firmware and bootloader, protecting against various attacks that attempt to inject malicious code or compromise the boot process.
- **Key rollback protection:**
  - To prevent attacks that involve rolling back to a previously signed firmware version with known vulnerabilities, Boot Guard includes protections against key rollback.
- **Configurability:**
  - System manufacturers have some flexibility in configuring Boot Guard based on their specific security requirements. They can, for example, decide which firmware and bootloader components are subject to verification.
- **Integration with secure boot:**
  - Intel Boot Guard works in conjunction with other security technologies, such as UEFI Secure Boot. Secure boot ensures that only signed and authenticated code is allowed to run during the boot process.
- **OEM customization:**
  - Original Equipment Manufacturers (OEMs) can customize the Boot Guard policies to align with their specific security needs, allowing them to adapt the technology to their hardware implementations.

It is important to note that, while Intel Boot Guard enhances system security, it is just one component of a comprehensive security strategy. Secure firmware, secure boot, and other security features collectively contribute to creating a more resilient and secure computing environment. Additionally, the specifics of Boot Guard may vary across different Intel processor generations, so it is advisable to refer to Intel's official documentation for the most accurate and up-to-date information.

---

## AMD Platform Secure Boot (PSB)

AMD Platform Secure Boot (PSB) is a security feature designed to enhance the security of AMD processors and platforms by focusing on the boot process. PSB is part of AMD's security initiatives to protect against unauthorized code execution during the system boot-up process.

Key features of AMD Platform Secure Boot include:

- **Secure boot mechanism:**
  - AMD PSB is a secure-boot mechanism that ensures the integrity of the boot process by verifying the authenticity of firmware and bootloader components before allowing them to execute.
- **Protection against unauthorized code execution:**
  - PSB helps protect the system from threats related to unauthorized or malicious code attempting to run during the boot sequence.
- **Integration with industry standards:**
  - AMD PSB is designed to work in conjunction with industry-standard secure boot protocols, such as UEFI Secure Boot.
- **Chain of trust:**
  - PSB establishes a chain of trust from the initial firmware load through the bootloader and into the operating system, ensuring that each step in the boot process is verified and secure.
- **Cryptographic verification:**
  - Cryptographic methods, such as digital signatures, are used to verify the authenticity and integrity of firmware and bootloader code. Only code with valid signatures is allowed to run.
- **Protection against rootkits and bootkits:**
  - By securing the boot process, PSB helps defend against certain types of attacks, including rootkits and bootkits, which aim to compromise the system at an early stage of boot-up.
- **OEM customization:**
  - Original Equipment Manufacturers (OEMs) can configure and customize AMD PSB settings based on their specific security requirements. This flexibility allows OEMs to adapt the security features to their hardware implementations.
- **Secure deployment of virtualization:**
  - In virtualized environments, PSB can contribute to the security of the hypervisor and virtual machines by ensuring a secure-boot process for the entire virtualization stack.

It is important to note that the specifics of AMD PSB and its integration with different processor generations may vary. For the most accurate and up-to-date information about AMD Platform Secure Boot, refer to AMD's official documentation. Security features are continually evolving, and AMD may introduce enhancements or updates to its security technologies over time.

---

## Security Protocol and Data Model (SPDM)

To defend against attacks targeting mutable components in Cisco UCS systems, the Security Protocol and Data Model (SPDM) specification enables a secure transport implementation that challenges a device to prove its identity and the correctness of its mutable component configuration. This feature is supported on Cisco UCS servers starting with Cisco UCS Manager Release 4.2(1d).

SPDM defines messages, data objects, and sequences for performing message exchanges between devices over a variety of transport and physical media. It orchestrates message exchanges between Baseboard Management Controllers (BMCs) and end-point devices over a Management Component Transport Protocol (MCTP). Message exchanges include authentication of hardware identities accessing the BMC. The SPDM enables access to low-level security capabilities and operations by specifying a managed level for device authentication, firmware measurement, and certificate management. Endpoint devices are challenged to provide authentication, and BMC authenticates the endpoints and only allows access for trusted entities.

The UCS Manager optionally allows uploads of external security certificates to BMC. A maximum of 40 SPDM certificates is allowed, including native internal certificates. Once the limit is reached, no more certificates can be uploaded. User-uploaded certificates can be deleted, but internal/default certificates cannot.

An SPDM security policy allows you to specify one of three security-level settings. Security can be set at one of the three levels listed below:

- **Full security:**
  - This is the highest MCTP security setting. When you select this setting, a fault is generated when any endpoint authentication failure or firmware measurement failure is detected. A fault will also be generated if any of the endpoints do not support either endpoint authentication or firmware measurements.
- **Partial security (default):**
  - When you select this setting, a fault is generated when any endpoint authentication failure or firmware measurement failure is detected. There will NOT be a fault generated when the endpoint does not support endpoint authentication or firmware measurements.
- **No security:**
  - When you select this setting, there will NOT be a fault generated for any failure (either endpoint measurement or firmware measurement failures).

You can also upload the content of one or more external/device certificates into BMC. Using an SPDM policy allows you to change or delete security certificates or settings as desired. Certificates can be deleted or replaced when no longer needed.

---

## Deployment and management at scale

Cisco UCS can be deployed in three different ways, depending on your needs, infrastructure, and preferences. Cisco Unified Fabric deployments can take place with UCSM or with Cisco Intersight. Both network-accessible UI interfaces are HTML-5-based and routinely vetted for web-application vulnerabilities in the CDSL process. Deployments without a fabric (that is, standalone) are handled using the baseboard management console.

### Cisco Intersight

Cisco Intersight is a Software-as-a-Service (SaaS) cloud-based infrastructure lifecycle management platform or an on-premises appliance-based device that delivers simplified deployment, monitoring, and support of Cisco Unified Computing System™ (Cisco UCS). Systems managed with Intersight run in Intersight Managed Mode (IMM). Cisco Intersight can be used for both first-time deployment and management of UCS components as well as post-UCSM deployment management through a multifactor claim process.

In Cisco Intersight, a server profile enables resource management by streamlining policy alignment, and server configuration. You can create server profiles using the server profile wizard, or you can import the configuration details of Cisco UCS C-Series servers in standalone mode and fabric-interconnect-attached servers in Intersight Managed Mode (IMM) directly from Cisco Integrated Management Controller (IMC). You can create server profiles using the server profile wizard to provision servers, create policies to ensure smooth deployment of servers, and eliminate failures that are caused by inconsistent configuration. The server profile wizard groups server policies into the following four categories to provide a quick summary view of the policies that are attached to a profile:

- **Compute policies:** BIOS, boot order, and virtual media.
- **Network policies:** Adapter configuration, iSCSI boot, LAN connectivity, and SAN connectivity policies.
  - The LAN connectivity policy allows you to create Ethernet network policy, Ethernet network control policy, Ethernet network group policy, Ethernet adapter policy, or Ethernet QoS policy. When you attach a LAN connectivity policy to a server profile, the addresses of the MAC address pool, or the static MAC address, are automatically assigned.
- **Storage policies:** SD card and storage policies.
- **Management policies:** Device connector, IPMI over LAN, LDAP, local user, network connectivity, SMTP, SNMP, SSH, serial over LAN, syslog, NTP, certificate management, and virtual KVM policies.

There are three methods of running IMM. These depend on the type of deployment the end user needs. For example, if there is a requirement for an air-gapped environment, but IMM is needed, an on-premises version of Intersight can be used. Here are the types of Intersight deployments:

- Cloud-based Cisco UCS Management (Intersight SaaS)
  - This cloud-native approach provides a centralized, web-based interface accessible from anywhere. It simplifies IT management by eliminating the need for on-premises hardware and software, making it an ideal choice for organizations seeking agility, scalability, and easy access to the latest features.
- Connected virtual appliance (Intersight CVA)
  - For businesses that prefer an on-premises solution while still benefiting from cloud connectivity, the connected virtual appliance is the answer. It offers the flexibility to run the Intersight virtual appliance within the data center while maintaining seamless connections to the Intersight cloud for updates and technical support.
- Private virtual appliance (Intersight PVA)

- If security and isolation are important factors, the private virtual appliance delivers an on-premises, air-gapped option. It operates in complete isolation from the Intersight cloud and the internet, ensuring that the infrastructure remains secure while still taking advantage of Intersight's management capabilities.

**Each of these offer the following benefits:**

- Unified management
  - Intersight consolidates the management of compute, network, storage, and hyperconverged infrastructure, simplifying the management of complex IT environments.
- Automation and orchestration
  - Intersight automates routine processes, which, in turn, reduces errors and accelerates deployments.
- Optimized operations
  - With proactive monitoring and intelligent analytics, Intersight helps identify and resolve issues before they impact an organization's business, ensuring maximum uptime.
- Security and compliance (including hardware compatibility list [HCL] features)
  - Stay ahead of security threats with real-time security alerts and compliance checks, keeping infrastructure protected and compliant with industry standards. Cisco Intersight also offers comprehensive Hardware Compatibility List (HCL) features, ensuring that an organization's hardware is not only compatible but also fully supported for seamless operations. The HCL features provide detailed insights into hardware compatibility, allowing informed decision making and optimal infrastructure for peak performance.
- Intersight Monitoring Services (IMS)
  - IMS provides historical and real-time visibility of resource consumption and inventory health, policies for notifications based on thresholds and anomaly detection, and actions and recommendations for automated infrastructure changes in response to events. IMS helps reduce operational costs, prevent SLA violations, and increase system reliability and availability.

**Intersight Service Level Objectives and Agreements**

Intersight Service Level Objectives (SLOs) are documented in the Cisco Trust Portal here:

<https://trustportal.cisco.com/c/r/ctp/trust-portal.html#/19645092958710743>

The SLO details what you can expect from Cisco in terms of handling cloud-based service outages and maintenance requirements. It explains what qualifies for these labels and what you can expect from Cisco in these events.

Note that the SLO from Cisco is distinct from the SLA and SLO from AWS, which provides the infrastructure for the cloud services, by contract, with Cisco. These policies are described below and are what Cisco, and by extension, the customer, can expect from the infrastructure as a whole.

The following two links show what AWS delivers for cloud providers such as Cisco:

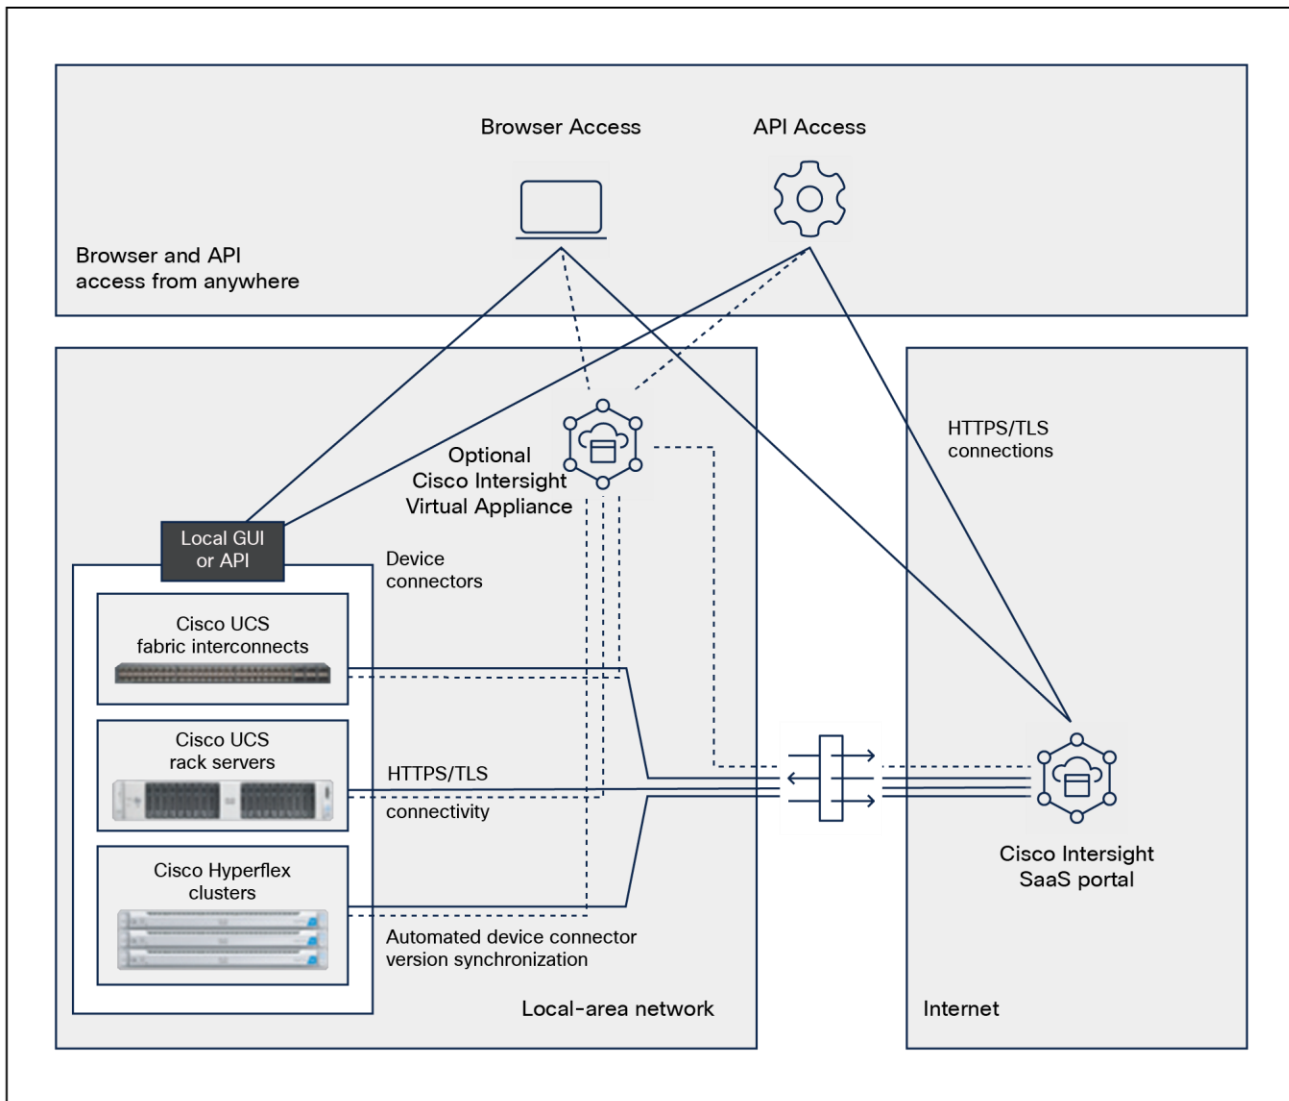
- AWS Service Level Agreements: [https://aws.amazon.com/legal/service-level-agreements/?aws-sla-cards.sort-by=item.additionalFields.serviceNameLower&aws-sla-cards.sort-order=asc&awsf.tech-category-filter=\\*all](https://aws.amazon.com/legal/service-level-agreements/?aws-sla-cards.sort-by=item.additionalFields.serviceNameLower&aws-sla-cards.sort-order=asc&awsf.tech-category-filter=*all)
- Reliability Pillar - AWS Well-Architected Framework: <https://docs.aws.amazon.com/wellarchitected/latest/reliability-pillar/welcome.html>

## Cisco Intersight Assist

Cisco Intersight Assist is a virtual machine that you can deploy in your data-center environment to streamline some Intersight device connection options. Intersight Assist helps you add endpoint devices to Cisco Intersight. A data center can have multiple devices that do not connect directly with Cisco Intersight. Any device that is supported by Cisco Intersight, but does not connect directly with it, will need a connection mechanism. Cisco Intersight Assist provides that connection mechanism and helps securely connect local data-center resources (VMware, storage, networking, etc.) to Cisco Intersight.

### Device connector

Cisco UCS systems are connected to the Intersight SaaS platform or on-premises virtual appliance through a device connector that is embedded in the management controller of each system.



**Figure 10.** Connection to Intersight services with the Cisco UCS device connector

---

The Intersight platform separates user and device traffic and communicates using industry-standard HTTPS and TLS protocols.

All data exchanged between devices and the Intersight platform uses industry-standard encryption and security protocols. Connected devices use Transport Layer Security (TLS) with restricted ciphers and HTTPS on the standard HTTPS port 443. All data sent to Intersight is encrypted using the Advanced Encryption Standard (AES) with a 256-bit, randomly generated key that is distributed with a public-key mechanism. In addition, every device connection to the portal is authenticated with a cryptographic token so that only legitimate devices can be managed, thus closing a potential Trojan horse attack vector. All connections are initiated from the device. Thus, firewalls can block all incoming connection requests; only HTTPS port 443 needs to be enabled for outbound connections. As a result, firewalls do not need any other special configuration to enable Intersight connectivity. Devices can be configured to use HTTPS proxy servers to add an additional layer of security through indirection.

To help ensure connection security and prevent man-in-the-middle attacks, Cisco UCS devices connecting directly to the Intersight platform use a single-destination HTTPS URL. The platform presents a certificate signed by a Certificate Authority (CA). If an unsigned certificate is presented, the devices will not connect to the portal. Intersight software and the device connector create a secure management framework that provides real-time information related to device security. This approach also allows connected devices and Intersight software to stay synchronized with the latest connection-security updates.

To monitor and manage devices with the Intersight platform, they first must be claimed from an Intersight account. Devices can be claimed using a browser by going to the SaaS or virtual appliance portal and clicking on the Claim Devices tab. Device IDs and a claim code, both of which are unique to the device, are retrieved from the device.

You can find the device ID and claim code through the device's local management interface. The claim code is refreshed every 10 minutes as an additional safeguard to ensure that the administrator claiming the device has physical access to it. Two-factor authentication is used to verify the identity and authenticity of each device being claimed. This authentication mechanism adds another layer of security to the device-claiming process. It requires access to the device as well as device identification information that is validated against your Intersight account.

In the event that an unauthorized user guesses or learns device information, the user cannot claim a device without physical access to the device.

### **Cisco UCS Manager**

Cisco UCS Manager (UCSM) enables you to manage general and complex server deployments. Systems deployed and managed with UCSM are in UCSM Managed Mode (UMM). For example, you can manage a general deployment with a pair of Fabric Interconnects (FIs), which are the redundant server access layer that you get with the first chassis or rack mount. This system can scale dramatically (see your hardware spec sheet for the deployment limits for your model). The system can be a combination of blades and rack mount servers to support the workload in your environment. As you add more servers, you can continue to perform server provisioning, device discovery, inventory, configuration, diagnostics, monitoring, fault detection, and auditing.

---

## Service profiles and policies in Cisco UCS

A service profile is a software definition of a server and its LAN and SAN connectivity. A service profile defines a single server and its storage and networking characteristics. Service profiles are stored in Cisco UCS fabric interconnects and are managed through specific versions of Cisco UCS Manager (which is the web interface for the fabric interconnect) or through purpose-written software using the API. When a service profile is deployed to a server, Cisco UCS Manager automatically configures the server, adapters, fabric extenders, and fabric interconnects to match the configuration specified in the service profile. This automation of device configuration reduces the number of manual steps required to configure servers, Network Interface Cards (NICs), Host Bus Adapters (HBAs), and LAN and SAN switches.

### Through policy we get security enforcement

In Cisco UCS Manager (UCSM), the service profile is a central component that defines the compute, network, and storage characteristics for a server within the UCS infrastructure. Essentially, it abstracts the physical hardware configuration from the logical configuration, enabling rapid provisioning, mobility, and scalability within the data-center environment. It is critical to note that policies prevent systems from being tampered with physically because changes to local settings are blocked when a policy is applied. This makes drift management nearly moot since the application of a policy prevents drift in the first place.

A service profile contains:

- **Hardware identity:** This includes details such as WWPN (world-wide port name) and WWNN (world-wide node name) for Fibre Channel, MAC addresses for Ethernet, BIOS settings, and firmware versions.
- **Boot policies:** These define how the server boots, which operating system it uses, and where it boots from (local disk, SAN, LAN, etc.).
- **Host firmware package:** This specifies the firmware versions to be applied to the server's components, ensuring consistency and compliance.

UCSM provides various policies that help enforce a secure deployment posture within service profiles:

- **Unified port policies:** These policies define the configuration for Ethernet and Fibre Channel ports, enabling administrators to set specific security-related parameters such as VLAN settings, port channel settings, QoS (Quality of Service) policies, etc.
- **Server BIOS policies:** These allow administrators to configure security-related settings in the server's BIOS, such as enabling or disabling specific hardware features, setting passwords, enabling secure boot, or configuring Trusted Platform Module (TPM) settings.
- **Local disk configuration policies:** These govern how local disks are configured, encrypted, or formatted, providing security measures for data stored on local disks.
- **Boot security policies:** These control boot-related security settings, including secure-boot configurations and control over boot devices.
- **Maintenance policies:** These define maintenance windows and policies that can help enforce security-related updates or configurations during specified maintenance periods.
- **Role-Based Access Control (RBAC):** This allows administrators to define roles and privileges, ensuring that only authorized personnel have access to critical UCSM functions and configurations, enhancing security by limiting access based on job responsibilities.

By utilizing these policies within UCSM when creating and managing service profiles, organizations can establish a more secure deployment posture, ensuring that servers are provisioned and configured according to

predefined security best practices and policies. This helps in reducing potential vulnerabilities and maintaining a standardized and secure computing environment within the Cisco UCS infrastructure.

### Cisco Integrated Management Console

Cisco UCS servers in standalone mode are managed using the baseboard management console, also called the Cisco Integrated Management Console (CIMC). The Cisco UCS C-Series rack server ships with the Cisco IMC firmware.

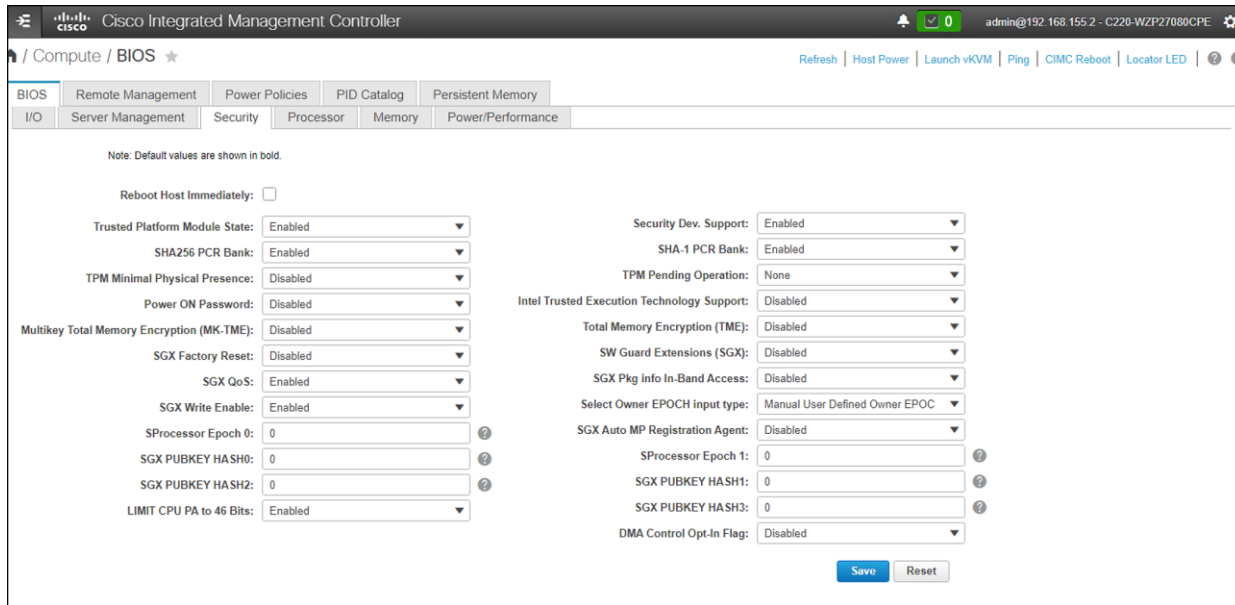
### Standalone deployment

Cisco IMC is a separate management module built into the motherboard. A dedicated ARM-based processor, separate from the main server CPU, runs the Cisco IMC firmware. The system ships with a running version of the Cisco IMC firmware. You can update the Cisco IMC firmware, but no initial installation is needed.

You can use a web-based GUI or SSH-based CLI or an XML-based API to access, configure, administer, and monitor the server. Almost all tasks can be performed in either interface, and the results of tasks performed in one interface are displayed in the other interface(s).

### CIMC security configuration

UCSM and Intersight both use policies and profiles to manage server groups. CIMC has these settings built-in individually on each server as part of the management firmware. Figure 11 shows a typical configuration screen for CIMC security settings.



**Figure 11.**  
Cisco IMC security configuration settings

---

## Default passwords

The Password Strength option is enabled by default on all management modes. Strong passwords must meet the following requirements:

- Must contain a minimum of 8 characters and a maximum of 64 characters
- Must contain at least three of the following:
  - Lowercase letters
  - Uppercase letters
  - Numbers
  - Special characters
- Must not contain a character that is repeated more than three times consecutively (for example, a password containing “aaabb” would be acceptable, but “aaaabbbb” would not).
- Must not be identical to the username or the reverse of the username
- Must pass a password dictionary check. For example, the password must not be based on a standard dictionary word.
- Must not contain the following symbols: \$ (dollar sign), ? (question mark), or = (equals sign)
- Should not be blank for local user and admin accounts

Additional password profile options:

- **Change count:** maximum times a password can be changed within the change interval.
- **Change interval:** time frame used by the change count.
- **No change interval:** minimum hours a local user must wait before changing newly created password.
- **Change during interval:** capability to change the password during the change interval.

After deployment and initial configuration are complete, make sure that any default passwords are changed or updated.

## Multifactor Authentication (MFA)

Cisco UCS Manager supports two-factor authentication for remote user logins. Two-factor authentication login requires a username, a token, and a password combination in the password field.

Two-factor authentication is supported when you use Remote Authentication Dial-In User Service (RADIUS) or Terminal Access Controller Access Control System Plus (TACACS+) provider groups with designated authentication domains with two-factor authentication for those domains.

With the implementation of Duo, the multifactor authentication is performed through a Duo-authentication proxy, which is an on-premises software service that receives authentication requests from your local devices and applications through RADIUS or LDAP; it optionally performs primary authentication against an LDAP directory or RADIUS authentication server and then contacts Duo to perform secondary authentication. Once the user approves the two-factor request, which is received as a push notification from Duo Mobile, or as a phone call or other notification, the Duo proxy returns access approval to the device or application that requested authentication.

---

## Access methods to management and configuration Interfaces

The management plane consists of functions that achieve the management goals of the system. Any management function undertaken by the user must rely on interaction through secure protocols, whether they are managing through a command line or a UI. This is handled through HTTPS for any UI access, whether UCSM, CIMC, or SaaS-based Intersight. Authenticated, tokenized access is used for in-house development through API. SSH for encrypted command line access is also supported. Management security also entails role-based access control as well as auditing and logging of system activities and user input, all of which are incorporated into every management mechanism.

Interactive management sessions using the command line take advantage of SSH or SCP. Either of these are available for UCSM or standalone CIMC deployments. These sessions take place in the embedded and abstracted management shell. This shell is hardened, does not allow root access, and cannot run user-space applications.

### Role-based access control

Local authentication is enabled by default. Use HTTPS and SSH for maximum security when accessing the Cisco UCS device. Numerous authentication methods provide enhanced security. There is a maximum of 48 local user accounts for RBAC access. Remote authentication uses LDAP, RADIUS, and TACACS+, with a maximum of 16 TACACS+ servers, 16 RADIUS servers, and 16 LDAP providers for a total of 48 providers. Roles defined in these domains are used to restrict and define access for different users. Refer to the deployment and configuration guides for your specific management RBAC configurations (UCSM, Intersight, or local CIMC).

### Authentication domains

The default (local) authentication and the console authentication can utilize different providers. Furthermore, authentication grouping uses a maximum of 16 groups and a maximum 8 providers per group. The provider authentication ordering method provides flexibility on what providers to use and what backups will be in place. The default authentication ports are configurable.

Default roles include AAA, admin, facility-manager, network, operations, read-only, server-equipment, server-profile, server-security, and storage. Additionally, roles can be customized by creating new roles and assigning privileges. The locales are used to define one or more organizations a user is allowed to access.

### Single sign-on with Intersight

Single Sign-On (SSO) authentication enables you to use a single set of credentials to log in to multiple applications. With SSO, you can log in to Intersight with your corporate credentials instead of your Cisco ID. Intersight supports SSO through SAML 2.0, acts as a Service Provider (SP), and enables integration with Identity Providers (IdPs) for SSO authentication. User authentication and role-based access control Intersight accounts form the authentication domain for users. The accounts control all resource access, and authenticated users are restricted from seeing any data in accounts where they are not authorized. With the SaaS platform, Cisco login IDs can be used for authentication with the identity provider for Cisco.com, which includes support for multifactor authentication. Both SaaS and on-premises Intersight implementations allow integration with external identity management systems to meet existing customer authentication requirements.

---

The Cisco Intersight framework uses granular access control with privileges managed per resource. Intersight software allows configuration of users and groups into several roles, and each user or group can be a member of multiple roles. Roles implemented include the following privileges:

- **Account administrator:** Full control and management capabilities for the Cisco Intersight account and devices under management.
- **Read-only:** Read-only visibility to resources under management.
- **Device technician:** Administrative device actions including device claim to a Cisco Intersight account.
- **Device administrator:** Administrative device actions including device delete from a Cisco Intersight account.
- **Server administrator:** Server lifecycle and policy-based management.
- **User access administrator:** User, group, and identity-provider configuration.

## **SSL key management: UI certificates and self-encrypting drives**

Cisco UCS ships with a self-signed certificate using a default 1024-length key pair. To employ a more secure method, use third-party certificates from a trusted source that affirms the identity of the Cisco UCS device.

Key management is also a core function for Self-Encrypting Drives (SEDs). SED keys can be managed either locally or remotely with a third-party key management server such as CipherTrust. Local key management requires a security key (passphrase) to be entered into the system. Remote key management requires configuration of the Key Management Server (KMS) and the proper distribution of certificates and public and private keys.

### **Cisco UCS IMC REST-based API**

Representational State Transfer (REST) or RESTful web services allow you to provide interoperability between systems on the internet. Using REST-compliant web services, you can access and manipulate web resources using a uniform and predefined set of stateless operations. Cisco has developed REST API capabilities to configure Cisco UCS C-Series servers using Redfish technology.

Redfish is an open industry standard specification and schema that specifies a RESTful interface and utilizes JSON and OData to help customers integrate solutions within their existing tool chains. Redfish is sponsored and controlled by the Distributed Management Task Force, Inc. (DMTF), a peer-review standards body recognized throughout the industry.

### **UCSM XML-based API**

The Cisco UCS Manager XML API is a programmatic interface to Cisco UCS. The API accepts XML documents through HTTP or HTTPS. Developers can use any programming language to generate XML documents that contain the API methods. Configuration and state information for Cisco UCS is stored in a hierarchical tree structure known as the management information tree, which is completely accessible through the XML API.

The Cisco UCS Manager XML API supports operations on a single object or an object hierarchy. An API call can initiate changes to attributes of one or more objects such as chassis, blades, adapters, policies, and other configurable components.

---

Authentication methods authenticate and maintain the session. For example:

- **aaaLogin:** Initial method for logging in
- **aaaRefresh:** Refreshes the current authentication cookie
- **aaaLogout:** Exits the current session and deactivates the corresponding authentication cookie

Use the aaaLogin method to get a valid cookie. Use aaaRefresh to maintain the session and keep the cookie active. Use the aaaLogout method to terminate the session (this also invalidates the cookie). A maximum of 256 sessions to the Cisco UCS can be opened at any one time.

### Monitoring

Server-system monitoring is crucial for security because it provides real-time visibility into the performance, health, and activities of servers within an IT infrastructure. Monitoring helps identify and respond to potential security threats, vulnerabilities, and irregularities, contributing to a proactive and effective security posture. Effective monitoring provides the following:

- Early detection of anomalies
  - Detect abnormal patterns or behaviors on servers, which may indicate a security incident. Early detection allows for a quicker response to potential threats.
- Identification of security incidents
  - Identify security incidents such as unauthorized access, malware infections, or suspicious network activities
- Visibility into system health
  - Insights into the overall health and performance of servers. A sudden drop in performance or unexpected system behavior may indicate a security compromise or the presence of malicious activities.
- Alerts and notifications
  - Generate alerts and notifications when predefined thresholds or security policies are breached.
- Resource utilization monitoring
  - Unusual spikes in CPU, memory, or network usage may indicate a security incident, such as a Denial-of-Service (DoS) attack or a compromised system engaging in malicious activities.
- Log analysis for security events
  - Analyze server logs for security-related events. This includes authentication attempts, access logs, and error messages that may reveal signs of unauthorized access or other security incidents.

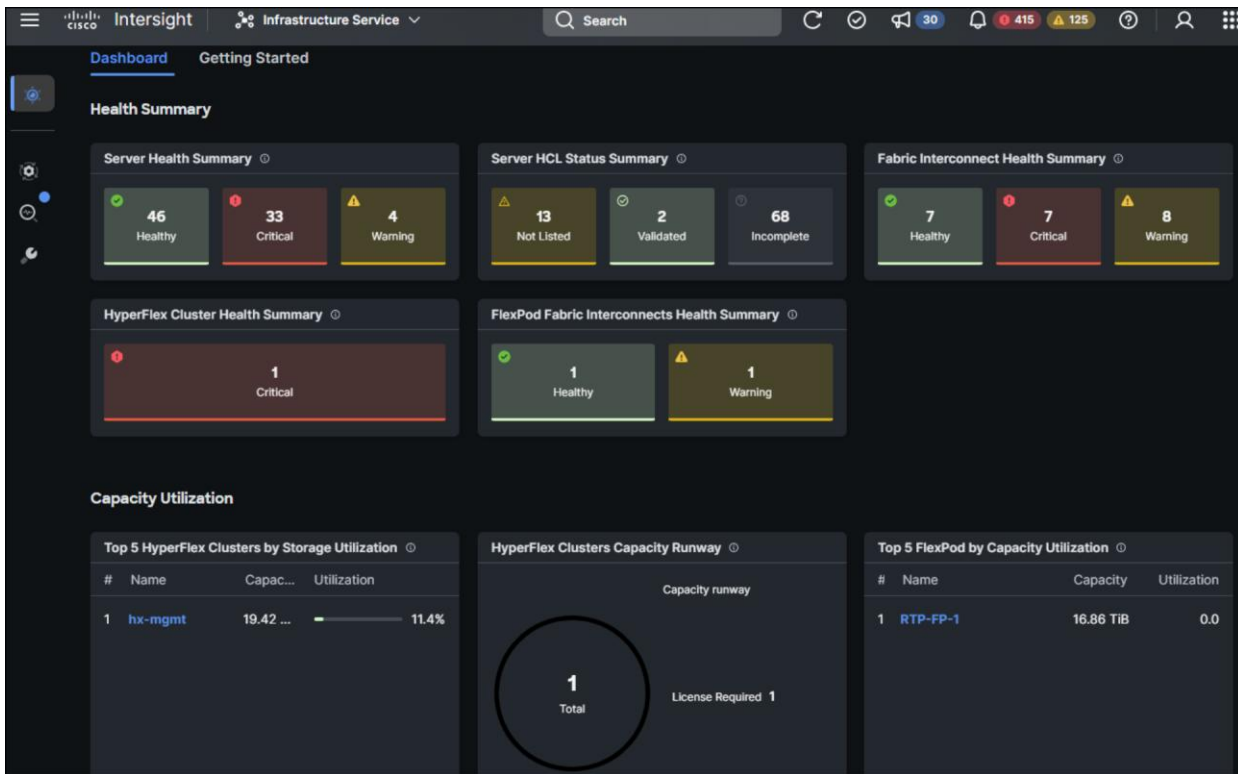
- 
- User activity monitoring
    - Monitoring user activities on servers helps in detecting suspicious behavior, such as unauthorized access or privileged users performing unexpected actions.
  - Compliance and auditing
    - Server monitoring helps provide the necessary data for audits. It verifies that security policies are enforced and that systems are in compliance with industry or organizational standards.
  - Incident response and forensics:
    - In the event of a security incident, monitoring data serves as valuable forensic evidence. It helps security teams understand the nature of the incident, trace its origins, and implement corrective measures to prevent future occurrences.
  - Patch and update management
    - Monitoring systems can track the status of server patching and updates. Ensuring that servers are up to date on security patches is essential for protecting against known vulnerabilities.
  - Capacity planning for security resilience
    - Monitoring assists in capacity planning, allowing organizations to anticipate resource demands and ensuring that servers are equipped to handle security-related loads, such as increased traffic during a Distributed Denial-of-Service (DDoS) attack.

By actively monitoring server systems, organizations can enhance their ability to prevent, detect, and respond to security threats effectively. It is a foundational element of a comprehensive security strategy, providing the insights needed to maintain a secure and resilient IT environment.

### **Monitoring with Intersight**

Cisco Intersight provides a dashboard for real-time health and inventory status monitoring. You can create, customize, rename, and manage multiple dashboard views by adding, removing, or rearranging widgets. In the Widget Library, you can select the widget(s) that you want to pin to the dashboard, preview the details for a server or a cluster, search for a widget, and add a custom title to a widget. You can toggle between the detail and list view of the available widgets in the library. You can add multiple instances of a widget to monitor different targets. You can add a maximum of 30 widgets per dashboard.

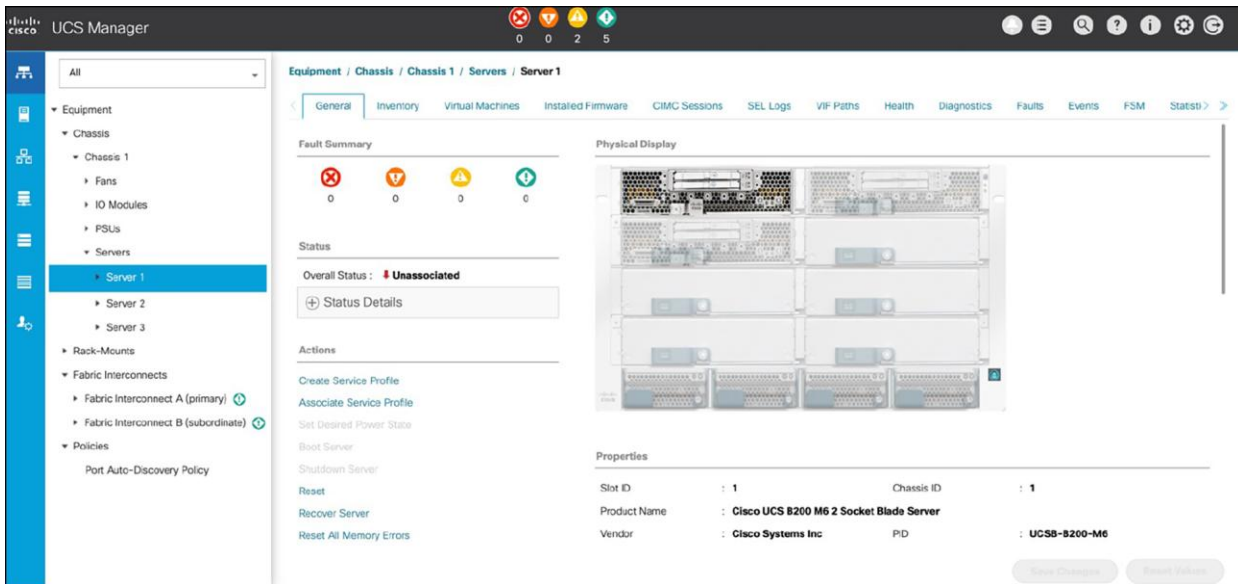
Intersight Monitoring Services (IMS) provide historical and real-time visibility of resource consumption and inventory health, policies for notifications based on thresholds and anomaly detection, and actions and recommendations for automated infrastructure changes in response to events. IMS help reduce operational costs, prevent SLA violations, and increase system reliability and availability.



**Figure 12.**  
Intersight monitoring dashboard

### Monitoring with UCSM

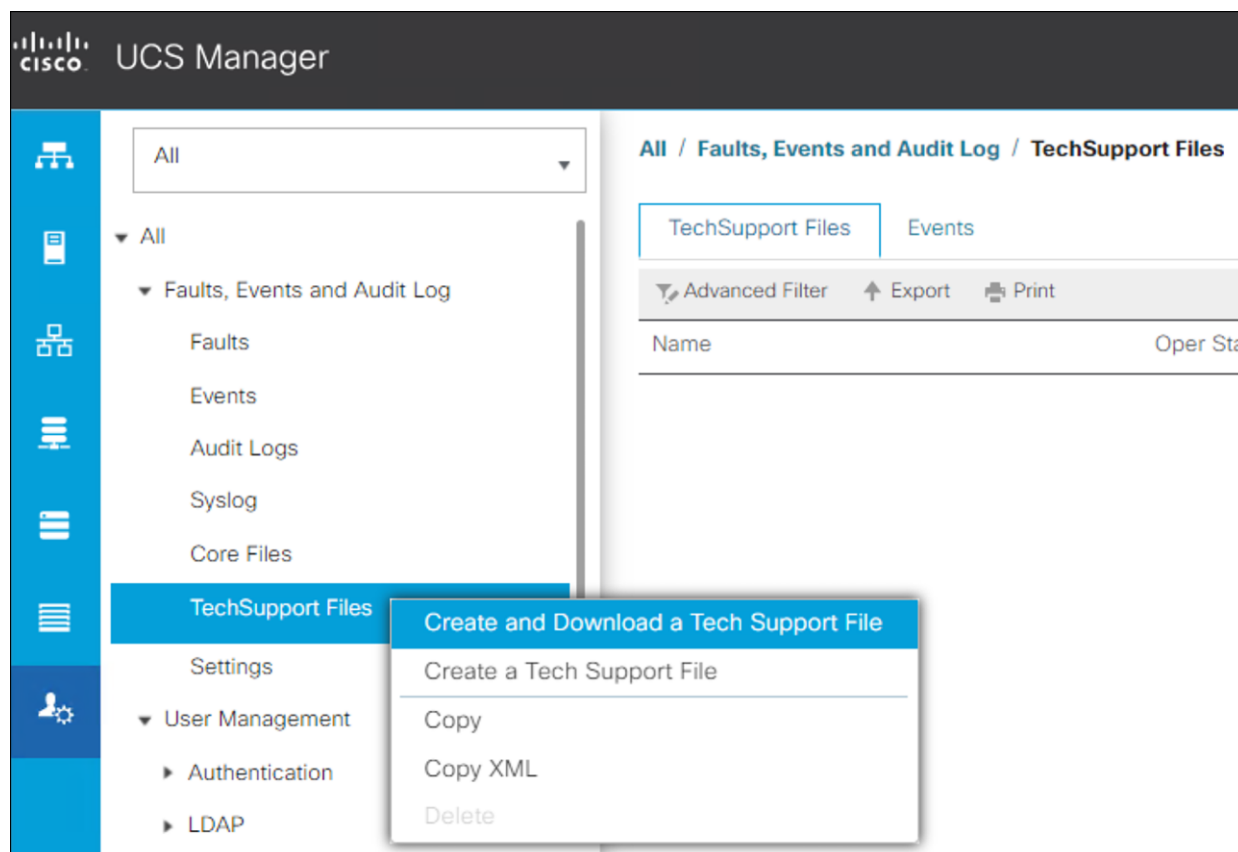
Cisco UCS Manager can detect system faults: critical, major, minor, and warnings. Cisco recommends that you monitor all faults of either critical or major severity status; immediate action is not required for minor faults and warnings.



**Figure 13.**  
UCSM fault and system summary

Logging is critical for investigation, audit, forensic analysis, and insight into overall system behavior and performance. The following logs are maintained by UCSM and can be examined at any time:

- System log.
  - System logs including faults, failures, and alarm thresholds (syslog).
  - The three types of syslogs: fault, event, and audit logs.
  - Global Fault Policy and settings that control syslogs.
- System event Log.
  - System hardware events for servers and chassis components and their internal components (System Event Log [SEL] logs).
  - The SEL policy that controls SEL logs.
- Simple Network Management Protocol (SNMP).
  - SNMP for monitoring devices from a central network management station and the host and user settings.
  - Fault-suppression policies for SNMP traps, call-home notifications, and specific devices.



**Figure 14.**  
Logging in UCSM

---

Monitoring with UCSM also includes:

- Statistics collection and threshold policies for adapters, chassis, host, ports, and servers.
- Call-home and Cisco Smart Call Home embedded device support.
- Hardware monitoring using the Cisco UCS Manager user interface.
- Cisco NetFlow Monitor for IP network traffic accounting, usage-based network billing, network planning, security, denial-of-service monitoring capabilities, and network monitoring.

### **Audit records**

To gain an understanding of existing, emerging, and historic events that are related to security incidents, an organization should have a unified strategy for event logging and correlation. This strategy must leverage logging from all network devices and use prepackaged and customizable correlation capabilities. Cisco UCS has a syslog capability that allows aggregation of logs at a centralized log server.

After centralized logging is implemented, a structured approach must be developed to analyze logs and track incidents. Based on the needs of the organization, this approach can range from a simple diligent review of log data to an advanced rule-based analysis.

The Cisco UCS audit log has a maximum of 10,000 entries. It utilizes a circular, FIFO logging design, so oldest entries will drop off as new entries are created. If you are at the 10,000 entry limit, you can rotate the logs currently in the audit log on your centralized server or export the log file as a csv from the GUI, thus clearing the audit log; you can also issue commands to pull the logs through (for example) PowerShell.

### **Decommissioning**

Securely decommissioning a server is a critical process to ensure that sensitive data is properly handled, and the server is retired in a way that minimizes the risk of data breaches or unauthorized access. Failing to decommission a server securely can lead to data exposure, legal and regulatory issues, and potential harm to an organization's reputation. Below are the levels of importance and the methods for securely decommissioning a server and its data.

Decommissioning a system or components of a system—specifically, the drives—requires special considerations in many circumstances. It is not sufficient to simply remove a drive or rotate a system out of production without sanitization. For older versions of UCS firmware, there are third-party applications that will run NIST-approved sanitization routines on plain text drives or encrypted drives. The Commission Regulation (EU) 2019/424 requires that data be securely disposed of. Secure data disposal is accomplished by using commonly available tools that erase the data from the various drives, memory, and storage in Cisco UCS servers and reset them to factory settings. You must be familiar with what devices are present in your UCS server and run the appropriate tools for secure data deletion. In some cases, you may need to run multiple tools.

Beginning early 2024, UCS firmware supports disk and system sanitization. This can more appropriately be termed data sanitization. Cisco IMC supports this NIST 800-88 compliant data sanitization feature. Using the data sanitization process, Cisco IMC erases all sensitive data, thus making extraction or recovery of user data impossible. As Cisco IMC progresses through the erase process, the status report is updated. You can check the status and progress of the data sanitization process for each individual device erase from the report. Cisco IMC reboots when the data sanitization process is completed and generates a report.

---

The erase process for data sanitization is performed in the following order on the server components:

- Storage
- VIC
- BIOS
- Cisco IMC

You can choose to either perform data sanitization on all the server components or select only VIC and Storage components for data sanitization.

Proper decommissioning reduces the risk of data breaches. If data is not securely wiped, deleted, or destroyed, it may be accessible to unauthorized individuals who can exploit it for malicious purposes.

Secure decommissioning is essential for compliance with data protection and privacy regulations. Many regulations, such as GDPR or HIPAA, require organizations to safeguard data, even during disposal. It also serves to protect an organization's intellectual property.

Additional procedural steps can be taken as well. These include physical destruction and environmentally responsible recycling of components where appropriate. If this is not possible because the systems will be reused, other things can be done such as disabling and removing all user accounts and resetting server configurations.

Secure decommissioning is a comprehensive process that involves technical, procedural, and organizational measures. By following these methods, organizations can minimize the risks associated with retiring servers and ensure that sensitive data is handled responsibly and securely.

#### **Instant Secure Erase (ISE) drives**

[Instant Secure Erase \(ISE\)](#) is a drive erasure method that provides a fast way to delete all data quickly when required. It encrypts the drive, and when it is time to permanently delete all data, only the encryption key has to be deleted. ISE is a super set of noncryptographic disk secure and it utilizes encryption to make data unreadable. It contains noncryptographic secure erase commands, but it also adds a "Crypto" Erase (CE) command, which can be utilized by both hard disks and solid-state drives if available. This method deletes the drive in a few seconds compared to a noncryptographic secure erase command, which writes over all the sectors and can take many hours. Only specific qualified ISE drives are available for use with Cisco UCS. See your Cisco representative for more information.



**Figure 15.**

ISE drives securely erase all content on a disk in an unrecoverable, cryptographic manner.

- ISE is a super set of noncryptographic disk secure erase commands and utilizes encryption to make data unreadable.
- ISE contains secure erase commands, but it also adds a "Crypto" Erase (CE) command.
- Each disk creates a key that is used to encrypt/decrypt data as it is written or read.
- When the "crypto" erase command is accessed, the key is destroyed and all data on the disk is unable to be read. This happens instantaneously.
- Like SED secure erase (see the section on SEDs below) ISE only takes a few milliseconds to make the disk unreadable.
- **While ISE uses cryptographic techniques to securely erase data, it does not offer data encryption to protect data at rest. SEDs are required for this, and they must be locked with a KEK solution.**

---

## 5. Secure application operation

The next pillar we will examine in the secure Cisco UCS posture environment is application operation. It is crucial that the system be able to ensure that applications can run in a fenced and protected compute regime. To this end, confidential computing is used, in hardware, to ensure isolated, encrypted, and protected execution.

### **Confidential computing**

Confidential computing is a cloud-computing technology that isolates sensitive data in a protected CPU enclave during processing. The contents of the enclave – the data being processed, and the techniques that are used to process it – are accessible only to authorized programming code and are invisible and unknowable to anything or anyone else, including the cloud provider.

A confidential computing secure enclave refers to a protected and isolated environment within a computing system where sensitive data or operations can be securely processed or stored. This secure enclave ensures that the data within it is protected from unauthorized access, even from other parts of the system or privileged software layers.

Cisco UCS implements all of the processor vendor confidential computing capabilities from AMD and Intel. This includes trusted execution environments through secure enclaves as well as the various flavors of memory encryption. Implementation of these features is processor-dependent and chosen at build or purchase time.

The concept of secure enclaves is primarily focused on maintaining the confidentiality, integrity, and privacy of sensitive information, especially when dealing with critical data or executing sensitive operations. These enclaves use hardware-based security mechanisms to create isolated and trusted spaces within the system's memory or processing units, offering a high level of protection against various types of attacks, including those attempting to access or manipulate the enclave's contents.

As company leaders rely more and more on public- and hybrid-cloud services, data privacy in the cloud is imperative. The primary goal of confidential computing is to provide greater assurance to leaders that their data in the cloud is protected and confidential, and to encourage them to move more of their sensitive data and computing workloads to public cloud services.

For years, cloud providers have offered encryption services to help protect data at rest (in storage and databases) and data in transit (moving over a network connection). Confidential computing eliminates the remaining data security vulnerability by protecting data in use—that is, during processing or runtime.

### **How confidential computing works**

Applications process data, and to do this, they interface with a computer's memory. Before an application can process (encrypted) data, it must go through decryption in memory. Because the data is, for a moment, unencrypted, it is left exposed. It can be accessed, encryption-free, right before, during, and right after it has been processed. This leaves it exposed to threats like memory dump attacks, which involve capturing and using Random Access Memory (RAM) put on a storage drive in the event of an unrecoverable error.

The attacker triggers this error as part of the attack, forcing the data to be exposed. Data is also exposed to root-user compromises, which occur when the wrong person gains access to admin privileges and can therefore access data before, during, and after it has been processed.

---

Confidential computing fixes this issue by using a hardware-based architecture referred to as a Trusted Execution Environment (TEE). This is a secure coprocessor inside a CPU. Embedded encryption keys are used to secure the TEE. To make sure the TEEs are only accessible to the application code authorized for it, the coprocessor uses attestation mechanisms that are embedded within. If the system comes under attack by malware or unauthorized code as it tries to access the encryption keys, the TEE will deny the attempt at access and cancel the computation.

This allows sensitive data to stay protected while in memory. When the application tells the TEE to decrypt it, the data is released for processing. While the data is decrypted and being processed by the computer, it is invisible to everything and everyone else. This includes the cloud provider, other computer resources, hypervisors, virtual machines, and even the operating system.

Intel and AMD offer different technologies and approaches to achieve confidential computing secure enclaves within their respective processor architectures. Below is a comparison between Intel's technologies (SGX, TDX, and TME) and AMD's features (SEV and SME) in terms of their approaches, functionalities, and key characteristics.

#### **Intel's technologies**

**Intel SGX** (Software Guard Extensions) creates isolated secure enclaves within the CPU's memory, allowing applications to protect sensitive code and data. This enables developers to create isolated execution environments for applications, protecting data and code even from higher-privileged software layers. It provides memory encryption, secure execution, remote attestation, and isolation.

**Intel TDX** (Total Memory Encryption and Intel Trusted Execution Technology) focuses on enhancing security in virtualized environments by providing memory encryption and secure execution environments for virtual machines. It provides total memory encryption, secure-boot processes, and hardware-based isolation to protect against attacks in virtualized environments. TDX protects VMs from unauthorized access and tampering, ensures secure migrations, and provides a trusted execution environment.

**Intel TME** (Total Memory Encryption) encrypts system memory to safeguard against unauthorized access, ensuring data confidentiality even if an attacker gains physical access to the memory. It protects system-memory contents through encryption, ensuring data confidentiality and integrity. TME aims to prevent data breaches and unauthorized access to memory contents.

#### **AMD's technologies**

**AMD SEV** (Secure Encrypted Virtualization) focuses on enhancing security in virtualized environments by providing hardware-based memory encryption for VMs. It offers memory encryption for each VM, isolating them from each other and the hypervisor, protecting against attacks in cloud environments. SEV provides memory encryption and isolation and facilitates secure VM migrations between physical hosts.

**AMD SME** (Secure Memory Encryption) encrypts the system's memory, protecting against unauthorized access and physical attacks by encrypting memory contents. It encrypts system-memory contents transparently without requiring specific software modifications, protecting against memory snooping attacks. SME protects memory contents through encryption, enhancing security against physical attacks.

Comparing the two, both Intel and AMD technologies aim to provide hardware-based security mechanisms to protect sensitive data and create secure enclaves. Intel SGX and AMD SEV focus on creating isolated execution environments for applications or virtual machines, whereas Intel TME and AMD SME concentrate on encrypting system memory to protect against unauthorized access.

---

Intel's TDX is more tailored for virtualized environments, offering features for VM security, while AMD's SEV is similarly focused on enhancing security in virtualized environments. Each technology has its unique characteristics, such as Intel SGX's focus on secure execution or AMD SEV's capabilities for secure VM migrations.

Overall, both Intel and AMD technologies contribute significantly to confidential computing by offering hardware-based security features, encryption mechanisms, and isolation to protect against various threats and attacks targeting sensitive data and applications. The choice between these technologies often depends on specific use cases, system requirements, and compatibility with the existing infrastructure.

### **Why use confidential computing?**

In summary, confidential computing is critically important in server operations for the following reasons:

- To protect sensitive data, even while in use—and to extend cloud computing benefits to sensitive workloads. When used together with data encryption at rest and in transit with exclusive control of keys, confidential computing eliminates the single largest barrier to moving sensitive or highly regulated data sets and application workloads from an inflexible, expensive on-premises IT infrastructure to a more flexible and modern public cloud platform.
- To protect intellectual property. Confidential computing is not just for data protection. The TEE can also be used to protect proprietary business logic, analytics functions, machine learning algorithms, or entire applications.
- To collaborate securely with partners on new cloud solutions. For example, one company's team can combine its sensitive data with another company's proprietary calculations to create new solutions – without either company sharing any data or intellectual property that it does not want to share.
- To eliminate concerns when choosing cloud providers. Confidential computing lets a company leader choose the cloud-computing services that best meet the organization's technical and business requirements, without worrying about storing and processing customer data, proprietary technology and other sensitive assets. This approach also helps alleviate any additional competitive concerns if the cloud provider also provides competing business services.
- To protect data processed at the edge. Edge computing is a distributed computing framework that brings enterprise applications closer to data sources such as IoT devices or local edge servers. When this framework is used as part of distributed cloud patterns, the data and application at edge nodes can be protected with confidential computing.

There may be some limitations to the use of these technologies depending on the operating system you install and the applications that you intend to use. For example, if you enable SGX on a VMware ESXi system, you will not be able to vMotion VMs. Please consult your operating system and application vendors before deciding to implement these features.

---

## The Confidential Computing Consortium

In 2019, a group of CPU manufacturers, cloud providers, and software companies – Alibaba, AMD, Baidu, Fortanix, Google, IBM/Red Hat, Intel, Microsoft, Oracle, Swisscom, Tencent, and VMware – formed the Confidential Computing Consortium (CCC) under the auspices of The Linux Foundation. Cisco is a member of the consortium.

The CCC's goals are to define industry-wide standards for confidential computing and to promote the development of open-source confidential computing tools. Two of the Consortium's first open-source projects, Open Enclave SDK and Red Hat Enarx, help developers build applications that run without modification across TEE platforms.

However, some of today's most widely used confidential computing technologies were introduced by member companies before the formation of the Consortium. For example, Intel SGX (Software Guard Extensions) technology, which enables TEEs on the Intel Xeon CPU platform, has been available since 2016; in 2018 IBM made confidential computing capabilities generally available with its IBM Cloud Hyper Protect Virtual Servers and IBM Cloud Data Shield products.

## 6. Secure data delivery and storage

The final pillar we will examine is the secure storage and delivery of data. This deals with encryption, key management, and data ingress/egress isolation. Traditional ciphers used in data encryption are nearing their functional end-of-life due to the encroaching capabilities of quantum computing. It is becoming increasingly important to consider and utilize post-quantum cryptography as part of a holistic approach securing data. To this end, Cisco announced membership in the Post-Quantum Cryptography Alliance in February 2024. The goal is to guide and implement quantum-resistant ciphers in the industry and across Cisco products.

### Encryption and key management

Encryption and remote key management play critical roles in ensuring secure data delivery, particularly in scenarios where sensitive information is transmitted or stored. These security measures contribute to protecting data confidentiality, integrity, and authenticity.

Encryption is primarily employed to ensure the confidentiality of data during transmission or while stored on a system. Cisco UCS has support for hardware-based encrypted drives (self-encrypting drives, or SEDs) and can maintain a local key or be configured to securely use a remote Key Management Server (KMS). This is encryption for data at rest (DARE). Data in transit can be encrypted in many ways, and UCS has a robust ecosystem to take advantage of all on-wire encryption solutions based on Cisco products. This can be accomplished in hardware (for example, on point-to-point or perimeter network devices) or by using “Cisco on Cisco” with the myriad virtual solutions that can run directly on a containerized UCS or hypervisor-based deployment. By encrypting data at both ends—during transmission and storage—organizations ensure that sensitive information remains secure throughout its lifecycle.

Key management is an important aspect of an encryption deployment. Remote key management involves securely storing encryption keys separate from the encrypted data. Keys are often considered as sensitive as the data they encrypt. By managing keys remotely and securely, organizations prevent a single point of failure and reduce the risk of unauthorized access to both data and keys.

Regularly rotating and updating encryption keys is a security best practice. Remote key management systems facilitate the secure rotation and distribution of new keys. This helps ensure that even if a key is compromised, the window of vulnerability is limited, and older keys are no longer in use.

## Encryption Key Management

### What is Key Management?

- Data is only as secure as the encryption keys
- **Key management** is the tasks involved with protecting, storing, backup and organizing keys
- Specialized vendors provide enterprise key management offerings
- Either IT or InfoSec teams manage the deployments

### UCS and Key Management?

- Local key and remote key management options
- UCS policy centric approach
- Enterprise key management integration
- Self-signed and CA signed certificates
- Workflows supported:
  - Disable/enable
  - Local key to remote key
  - Re-key
- KMIP 1.1 compliant

**Figure 16.**

UCS key management features

## Self-encrypting drives

Data-at-rest encryption on a Cisco UCS server can happen in software for VMs (for example, Vormetric Transparent Encryption in VMcrypt) or by the operating system (for example, Microsoft's BitLocker). This can also be accomplished using hardware. To that end, Self-Encrypting Drives (SEDs) were developed and have many advantages. SEDs have a negligible impact on performance speed and latency. The encryption process is completely integrated into the drive, so there is no need for other system components to step in and perform any heavy lifting. SEDs are independent of the operating system, so even if a hacker attacks a computer, it is nearly impossible to access the SED (and the encryption keys stored within) when the computer is turned off.

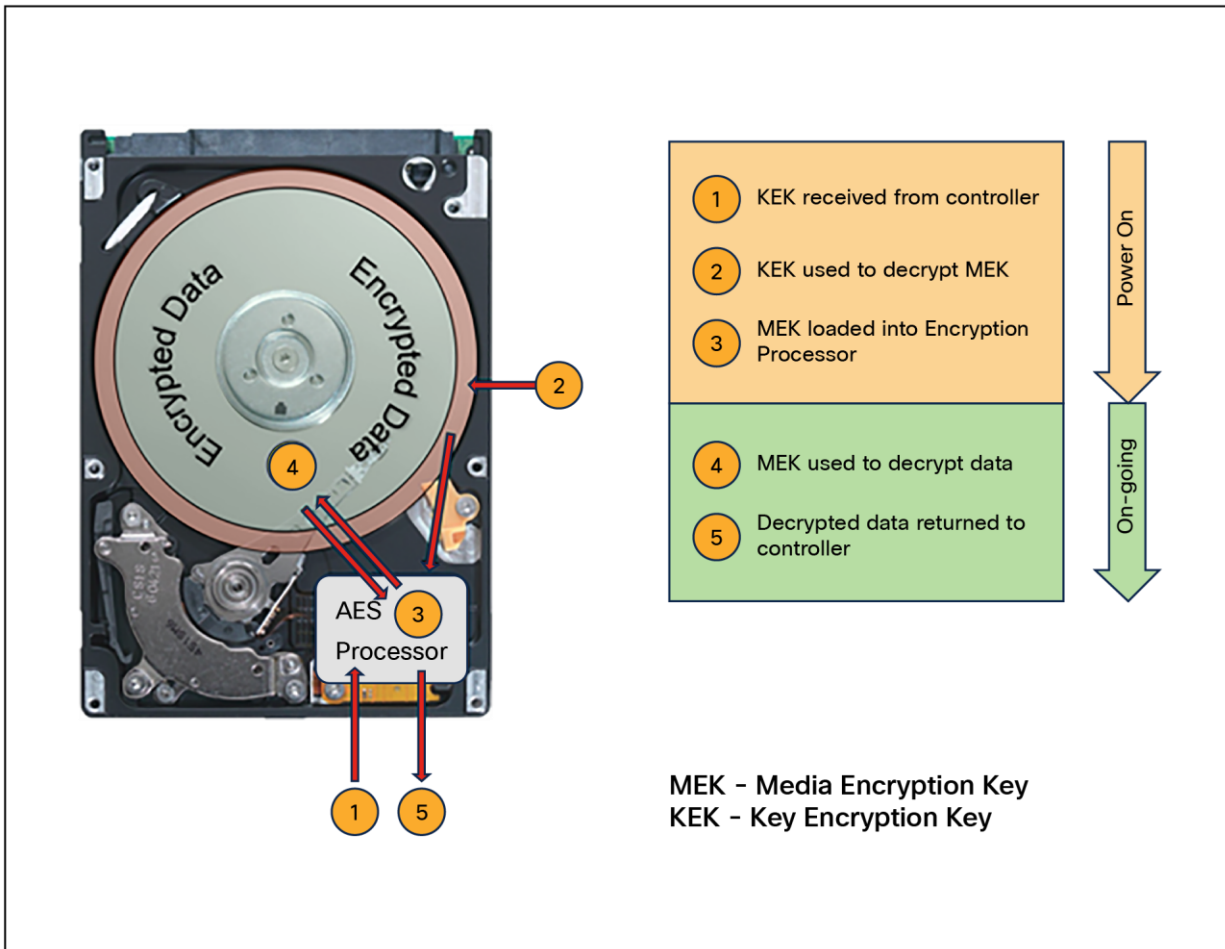
In a Cisco UCS server, SEDs can utilize a local key (security key) or a remote key management solution. Remote key management is the recommended method since it does not rely on stored or "remembered" pass phrases. The key management software optimizes the SED's decryption and encryption functions and key management, relieving the user of any active SED administration. Lastly, SEDs are inexpensive to deploy and maintain. SEDs encrypt the moment they come off the assembly line. Management software does the rest, ensuring that SEDs do their job without the need for human intervention.

Self-Encrypting Drives (SEDs) have special hardware that encrypts incoming data and decrypts outgoing data in real-time. The data on the disk is always encrypted in the disk and stored in the encrypted form. The encrypted data is always decrypted on the way out of the disk. A media encryption key controls this encryption and decryption. This key is never stored in the processor or memory. Both Cisco UCS Manager and Intersight support SED security policies on Cisco UCS C-Series servers, B-Series M5/6/7 servers, X-Series servers, and S-Series servers.

SEDs must be locked by providing a security key. The security key, which is also known as a key-encryption key or an authentication passphrase, is used to encrypt the media encryption key. If the disk is not locked, no key is required to fetch the data.

Cisco UCS Manager enables you to configure security keys locally or remotely. When you configure the key locally, you must remember the key. If you forget the key, it cannot be retrieved, and the data is lost. You can configure the key remotely by using a key management server (also known as KMIP server). This method addresses the issues related to safe-keeping and retrieval of the keys in the local management.

The encryption and decryption for SEDs is done through the hardware. Thus, it does not affect overall system performance. SEDs reduce disk-retirement and -redeployment costs through instantaneous cryptographic erasure. Cryptographic erasure is done by changing the media-encryption key. When the media-encryption key of a disk is changed, the data on the disk cannot be decrypted and is immediately rendered unusable. With Cisco UCS Manager Release 3.1(3), SEDs offer disk-theft protection for C-Series and S-Series servers.



**Figure 17.**  
Anatomy of a Self-Encrypting Drive (SED)

---

## Virtual interface card

A Cisco UCS VIC (virtual Interface card) plays a crucial role in enhancing secure data delivery within Cisco's Unified Computing System (UCS) infrastructure. The VIC is a key component that provides network connectivity for servers within the UCS platform.

The primary function of a Cisco UCS VIC is to provide network connectivity for servers. It supports Ethernet and Fibre Channel over Ethernet (FCoE) protocols, allowing servers to communicate with the network infrastructure securely. Cisco UCS utilizes a unified fabric approach, where both data and storage traffic share the same high-speed fabric. The VIC enables this convergence, simplifying cabling and providing a more efficient and flexible network architecture. UCS VICs support adaptive networking features, allowing them to dynamically adjust to changing network conditions. This adaptability helps optimize data delivery performance and responsiveness. The VIC also supports Quality-of-Service (QoS) features, allowing administrators to prioritize and manage network traffic based on application requirements. This helps ensure that critical data is delivered with the appropriate level of service.

Cisco UCS VICs are designed to work seamlessly with virtualized environments. They provide support for VMware, Microsoft Hyper-V, and other hypervisors, enabling secure data delivery in virtualized infrastructures. UCS servers, including the VIC, incorporate security features such as secure boot processes. The VIC itself supports secure boot; for example, the VIC firmware is signed, and the card contains a trust anchor that verifies its own validity upon boot. This helps establish a chain of trust during the server bootup, ensuring the integrity of the system and reducing the risk of compromise.

A VIC is the cornerstone of the UCS server's traffic isolation and segmentation when used in conjunction with a fabric interconnect. This can be essential for enhancing security by keeping different types of traffic separate, preventing unauthorized access to sensitive data.

By combining these features, a Cisco UCS VIC contributes to the secure and efficient delivery of data within the UCS infrastructure. It plays a central role in enabling unified fabric, supporting virtualized environments, and providing the necessary connectivity for secure and reliable data communication.

## Conclusion

This paper has provided an introduction to the Cisco UCS ecosystem and secure development process followed by an examination of various pillars of UCS system security.

Firstly, we looked at the chain of trust and how it relates to secure supply chain and manufacturing. This was followed by the system-level security features of Cisco UCS, which include secure system boot as well as policy-based deployments at scale. We looked at the three primary methods of managing the system(s) through UCSM, Intersight, or standalone CIMC. We wrapped this up with an introduction to the various monitoring methods available for each management mode.

Next, we examined what is involved with secure application operation, including hardware-based secure enclaves and confidential computing. Finally, we looked at secure data storage and delivery using encryption and how Cisco VICs aid in these efforts.

---

When we combine the inherent security features of the Cisco UCS platform with common-sense security practices such as the following:

- Maintenance of physical security
- Keeping server OS and firmware patched and updated to mitigate new threats
- Disabling functions that are not required
- Maintaining application security with RBAC, patching, and firewalls, and
- Storing and delivering data securely with encryption in hardware both on the server and on the wire through ecosystem design.

we can ensure that our server environments are as secure as possible.

## For more information

For additional information, see the following resources:

- [Cisco Trust Portal](#)
- [Cisco Security](#)
- [Cisco Security Advisories](#)
- [DMTF and Redfish](#)
- [Cisco UCS Manager XML API Programmer's Guide - Cisco UCS Manager XML API Method Descriptions \[Cisco UCS Manager\]](#)
- [UCSM multi-factor authentication](#)
- [Network guide \(ports, etc.\)](#)
- [File server security concerns](#)
- [Cisco Trustworthy Technologies](#)
- [Audit log entries and retention](#)
- [Cisco UCS Secure Data Deletion](#)
- [Cisco joins PQC consortium](#)
- [Post-Quantum Cryptography Alliance](#)
- [Industrial Design Security](#)

## Document information

Document summary	Prepared for	Prepared by
V1.0	Cisco Field	Aaron Kapacinkas
<b>Changes</b>		
N/A		

**Americas Headquarters**  
Cisco Systems, Inc.  
San Jose, CA

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

**Europe Headquarters**  
Cisco Systems International BV Amsterdam,  
The Netherlands

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

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: <https://www.cisco.com/go/trademarks>. Third-party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)