Intersight Managed Mode (IMM) Architecture Deep Dive

A journey under the hood

Vincent Esposito – Technical Solutions Architect
@vesposit

BRKCOM-3280

Session objectives

- Get in-depth knowledge about how Intersight Managed Mode (IMM) works
- Understand what happens under the hood during initial setup, deployment tasks & common operations
- Feel more confident operating an IMM infrastructure

Session non-objectives & warnings

- This session is not:
  - An introduction to UCS, Intersight or IMM
  - A deep dive session on Intersight Standalone Mode (ISM)
  - A troubleshooting session

- This is an "Advanced" (BRKCOM-3280) session
  - Expect some very technical content – slides flagged with 
  - It is there for you to better understand how the underlying architecture works & use as reference in the future

Webex App

Questions?

Use the Webex app to chat with the speaker after the session

How

1. Find this session in the Cisco Events mobile app
2. Click “Join the Discussion”
3. Install the Webex app or go directly to the Webex space
4. Enter messages/questions in the Webex space

Webex spaces will be moderated by the speaker until February 28, 2025.

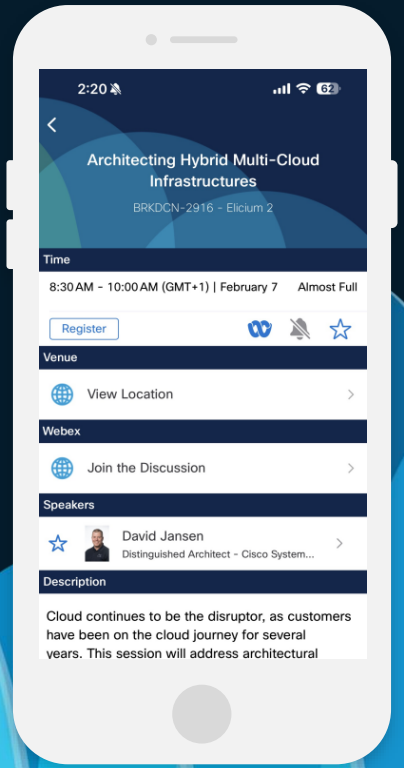

Agenda

- Introduction
- Intersight/IMM Architecture
- Setting up the UCS domain
- Discovering hardware
- Deploying Server & Chassis Profiles
- Performing (some) Operations
- Key Takeaways

Introduction

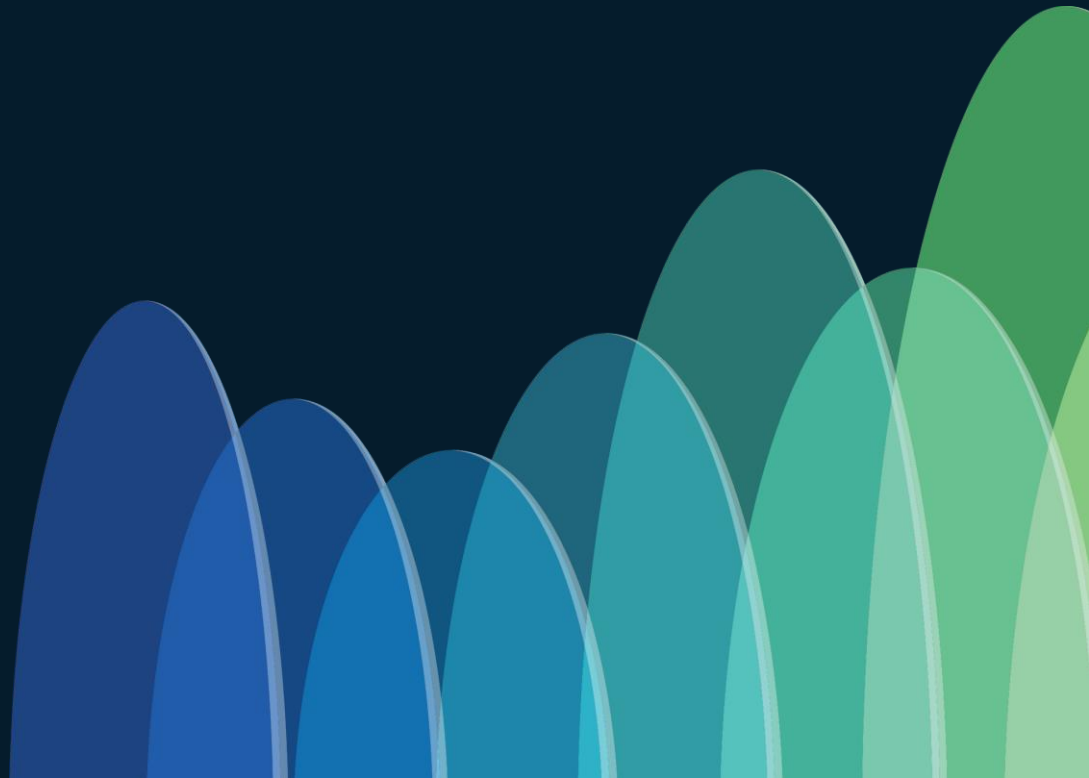

UCS Management with Intersight

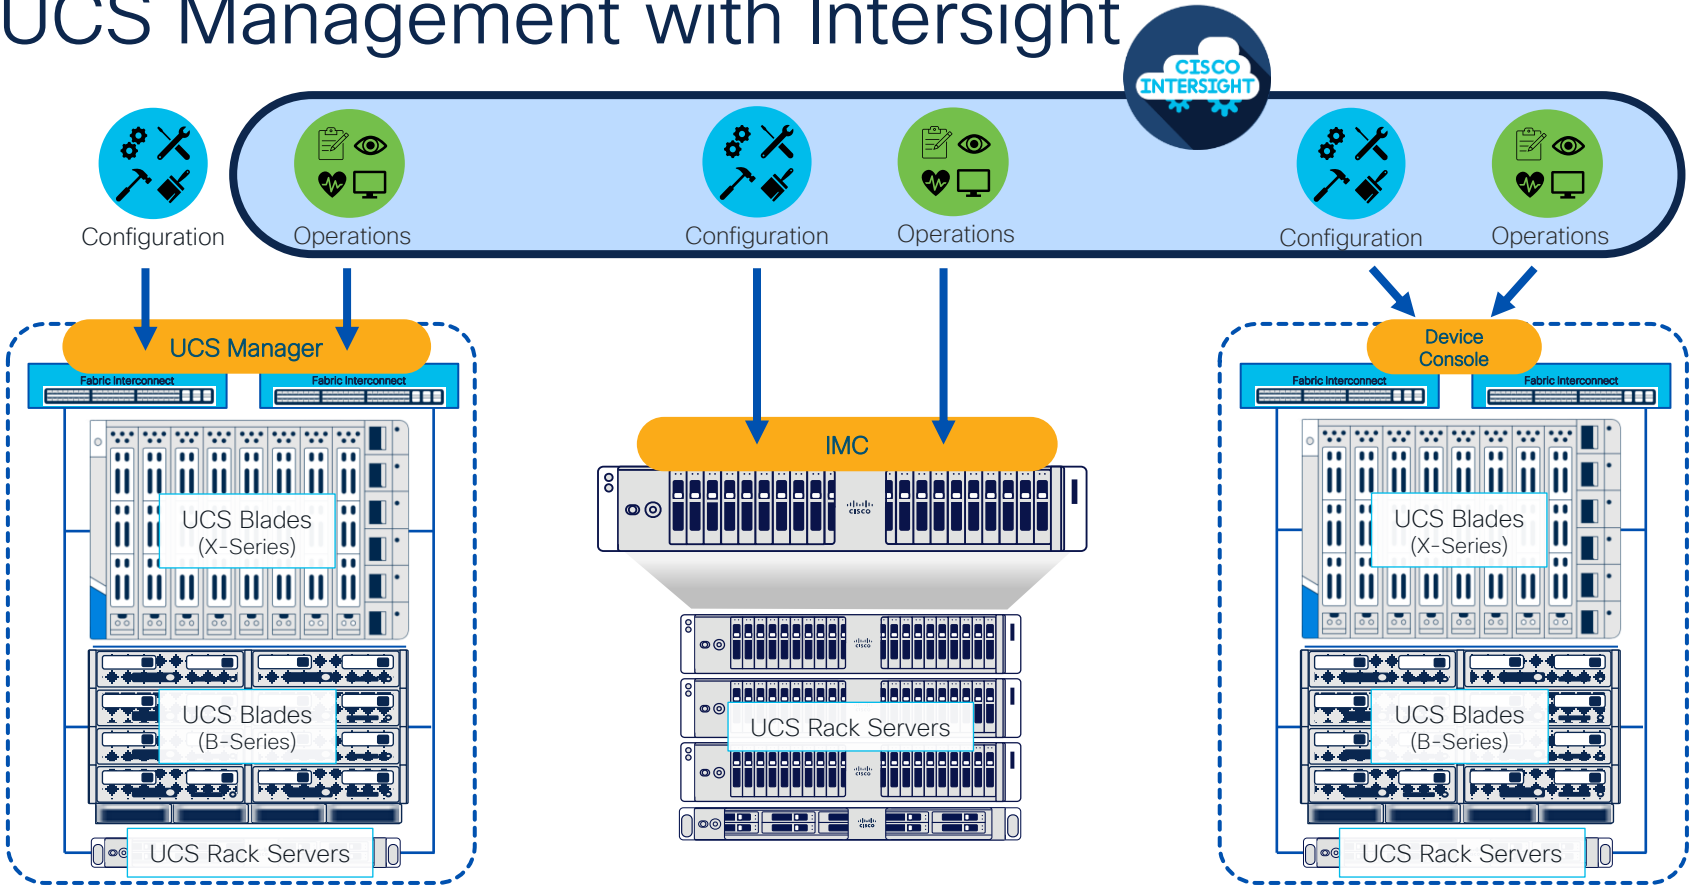

UCSM Managed Mode (UMM)

C-Series Standalone Mode (ISM)

Intersight Managed Mode (IMM)

Benefits of the IMM model

Modernize

- Feature velocity with CI/CD
- Modern RESTful API with OpenAPI
- **Redfish** standard

Simplify

- Common management framework for all Cisco Compute
- Merging of Local & Global Service Profile/Template (Multi domain)
- Policy Model simplification

Safer operations

- Configuration consistency for domains
- Decoupling of firmware / software features
- Improved Change Control with Assign/Deploy steps
- Better policy scale – not limited by FI resources

Intersight Managed Mode

Hardware & Software Requirements

Hardware

- FI: 6454/64108/6536/S9108-100G (x2)
- IOM: 2204/2208/2304/2408
- IFM: 9108-25G/9108-100G
- Servers: B/C/X-Series M5, M6, M7 & M8
- VIC: 1300/1400/15000

Firmware

- Infrastructure (A-bundle): 4.1(3) or higher
- Servers (B/C-bundles): 4.1(3) or higher
- 4.2(3)+ for most comprehensive support

Licensing

Intersight: Essentials or higher (all servers in domain need a license)

Scope

Entire UCS domain will be configured in Intersight Managed Mode (exclusive of UCS Manager)

Scale

20 chassis / 160 servers (same as UCSM)

https://intersight.com/help/saas/supported_systems#supported_hardware_for_intersight_managed_mode

Intersight Deployment Modes

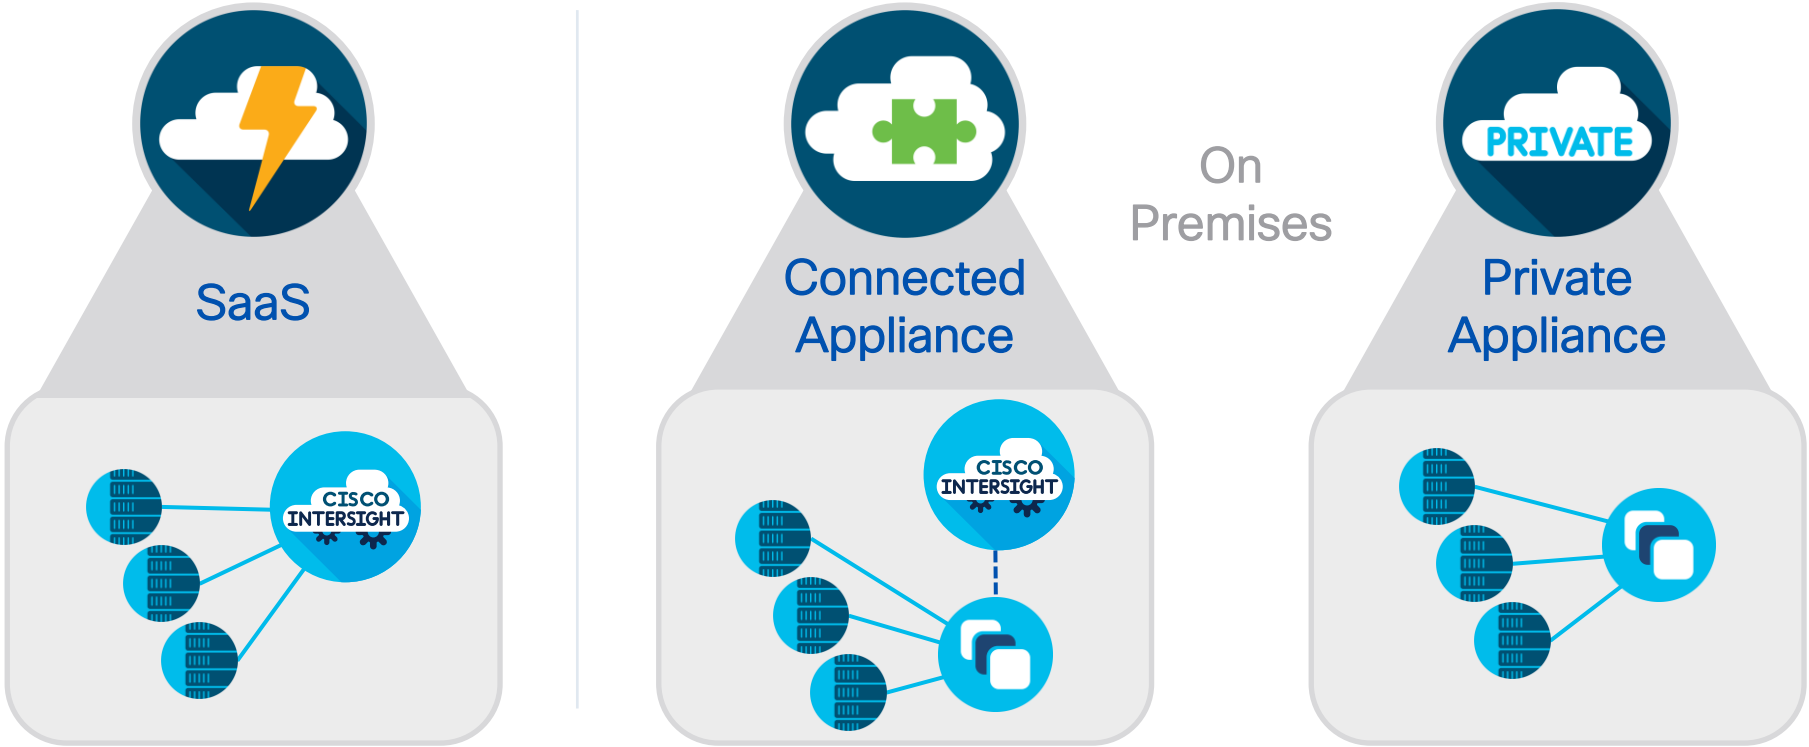

Intersight/IMM Architecture

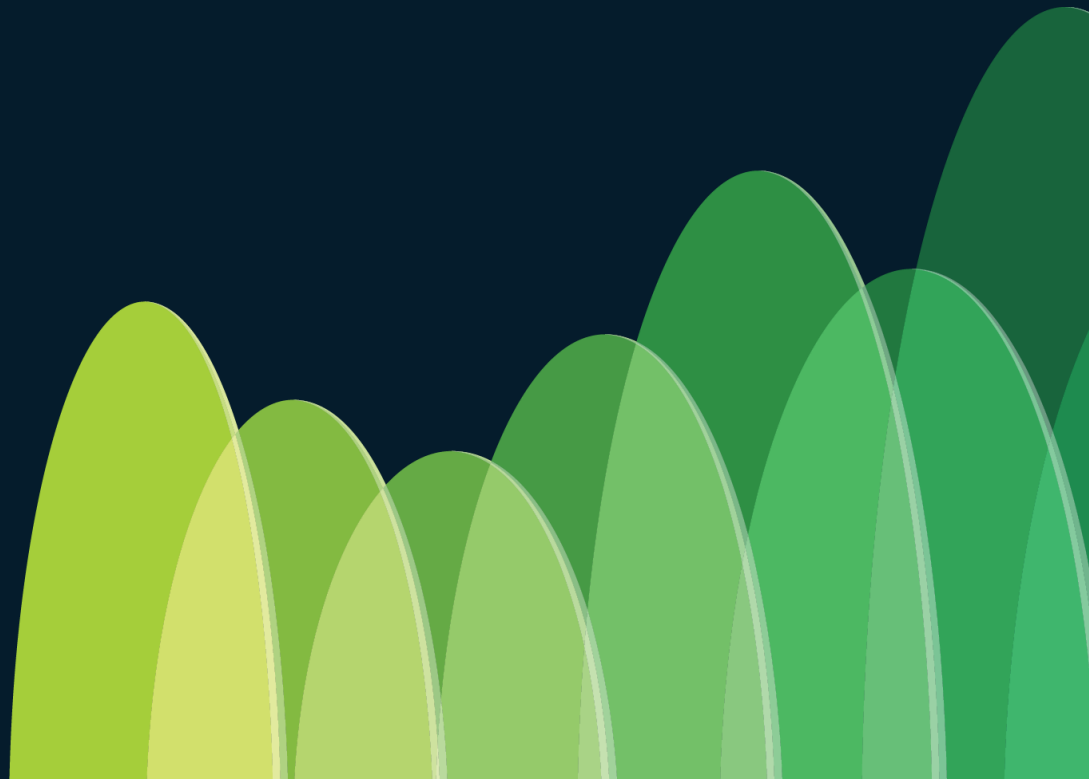

cisco Live!

Micro-services Architecture

- Intersight is fully built on micro-services
  - Continuous Delivery/Continuous Deployment (CI/CD)
  - Better horizontal scaling
- 100+ different micro-services spanning multiple product areas
  Platform Services, Server Management, Device Connector, Pools, Alarms, User Interface, Advisories, HCL, Cloud Orchestrator, etc.
- Intersight Appliance uses a lot of the same micro-services as SaaS + some specific to Appliance itself
- Uses Kubernetes for management, scaling & high-availability
- Leverages S3 service for storing firmware images / Tech-Support bundles / etc.
  - Local S3 service hosted internally for Appliance
  - Uses CDN for faster downloads on SaaS

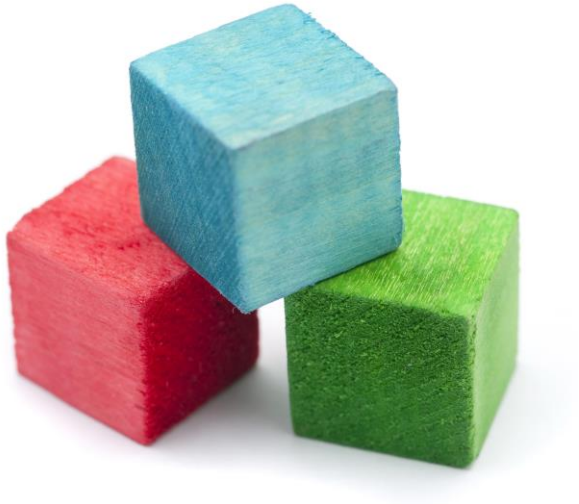

For your reference

Key Intersight micro-services

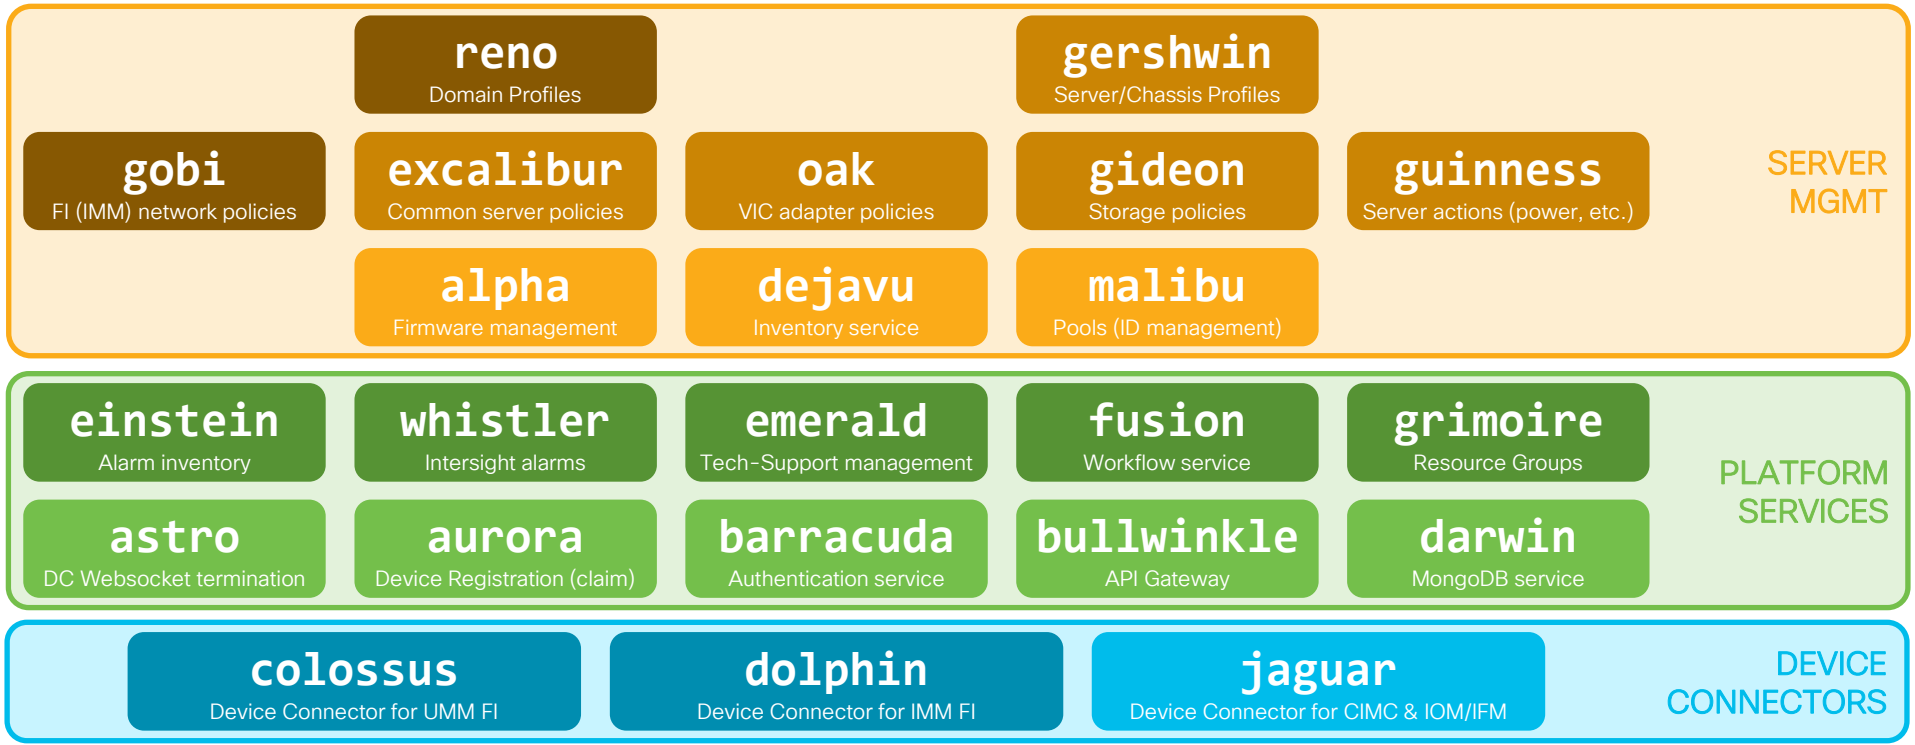

Key Intersight micro-services (Device Connector)

- **colossus:** Device Connector for FIs running in UMM
- **dolphin:** Device Connector for FIs running in IMM
- **jaguar:** Device Connector for CIMC (ISM & IMM) and IOM/IFM

Key Intersight micro-services (Platform Services)

- **astro:** Device Connector websocket termination (on Intersight side for SaaS)
- **aurora:** Device Registration (claiming)
- **barracuda:** Authentication service (manages IDPs, users & sessions)
- **bordeaux:** Front-end webserver (SaaS)
- **bullwinkle:** API gateway
- **darwin:** MongoDB service
- **einstein:** Alarms inventory
- **fusion:** Workflow service
- **grimoire:** Resource groups
- **whistler:** Intersight alarms

Key Intersight micro-services (Server mgmt)

- **alpha:** Firmware management service
- **dejavu:** Inventory of UCSM, IMM and ISM devices
- **dove:** Management package for FIs running in IMM (Device Console + CLI)
- **excalibur:** Policy service (for policies common to ISM and IMM)
- **gershwin:** Server/Chassis Profile management
- **gideon:** Storage Policy service
- **gobi:** FI (IMM) network Policy service
- **guinness:** Server actions (Power, LED, etc.)
- **magnum:** KVM sessions management
- **malibu:** Pools (ID management) service
- **oak:** VIC adapter Policy service
- **propeller:** Tunneled vKVM sessions
- **reno:** Domain Profile management
- **rio:** OS Install service

Key Intersight micro-services (Appliance)

- **equinox:** Device Connector for Intersight Appliance
- **echo**: Front-end webserver & Device Connector websocket termination
- **foster:** Vault service for storing credentials
- **fuji**: LDAP service
- **hammer:** Internal S3 service for storing images & Tech-Support bundles
- **hurricane:** Appliance admin settings
- **ketchup:** Appliance software management

Intersight Device Connector (DC)

A very light and autonomous piece of software allowing:

- Communication with the Intersight portal, wherever the portal is.
- Capable of inserting tasks / calls against infrastructure via pluggable / extensible framework

Key Features:

- Bundled with Firmware
- Embedded Product Feature
- Secure “Durable” Communications
- Self Updated
- Autonomous Check-In

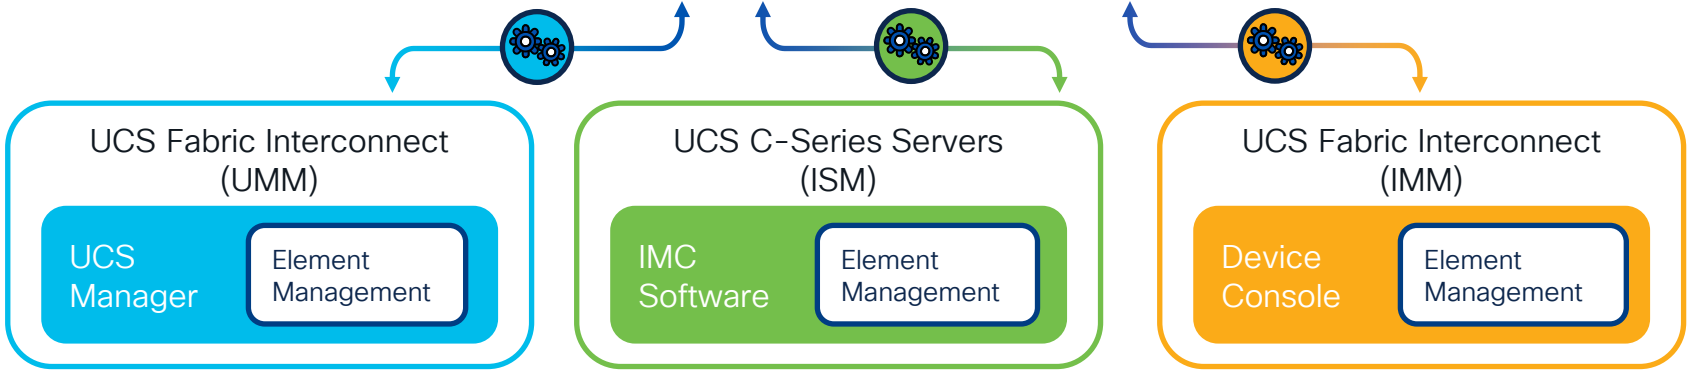

IMM Management Architecture Overview

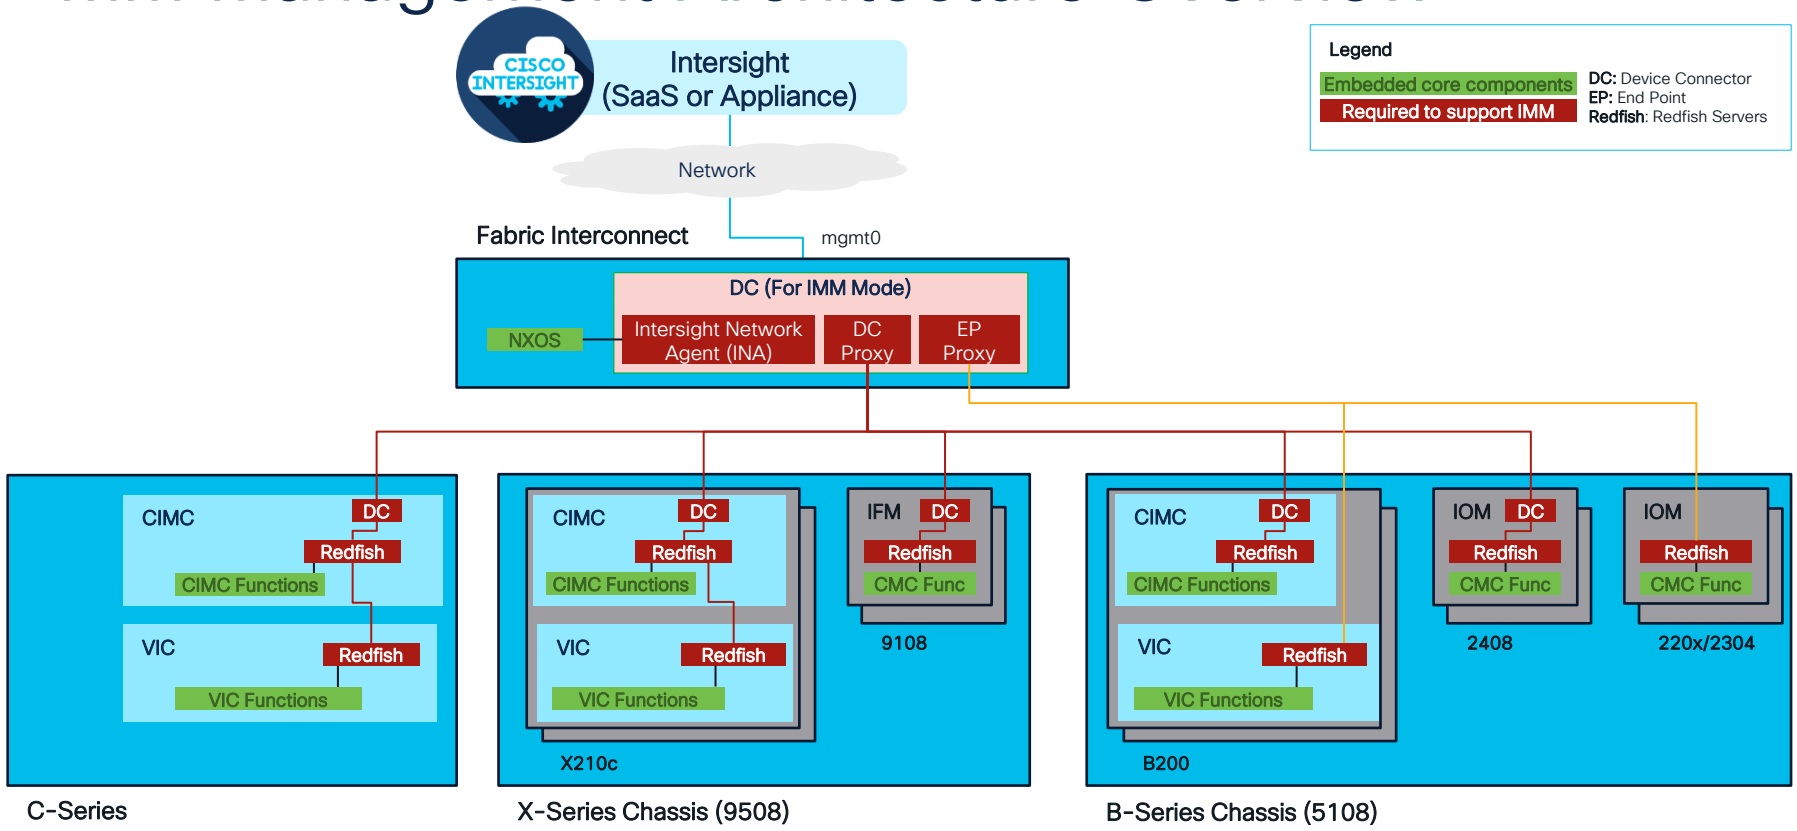

IMM Fabric Interconnect Management Architecture

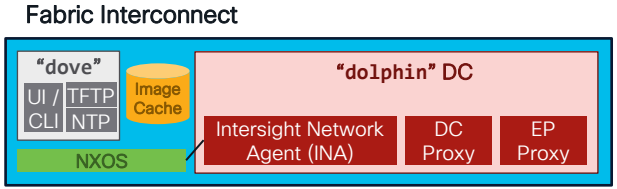

- Key Components:
  - Management Package ("dove") : Includes utilities to directly manage the UCS infrastructure (without Intersight) and required components like NTP (for time sync of all endpoints) and TFTP (for firmware upgrades)
  - Image Cache: Repository for caching firmware and utility images on the FI
  - NX-OS: Runs all switching aspects of the FI. Provides an NXAPI REST interface used by INA for management
  - Device Connector ("dolphin"): Enables FI to Intersight communication. Acts as parent DC to servers/chassis DC connected to the same IMM domain
  - Intersight Network Agent (INA): Responsible for all NXOS-related interactions including config and events
  - DC Proxy: provides connectivity to child DC of connected devices (CIMC, IOM/IFM)
  - EP Proxy: provides connectivity to endpoints that don't run a DC (IOM 220x, VIC)

IMM Server Management Architecture

- Key Components:
  - Device Connector (“jaguar”): Enables CIMC to Intersight communication. Acts as child DC to register automatically to Intersight via FI DC proxy
  - Redfish Server (CIMC): Provides Redfish API for managing CIMC
  - VIC Proxy: Proxy service for Redfish server running on VIC. Allows single entry point for all server components via CIMC DC. Only present on C-Series & X-Series
  - Redfish Server (VIC): Provides Redfish API for managing VIC. Communicates via VIC Proxy for C-Series & X-Series. Communicates via FI EP Proxy for B-Series.

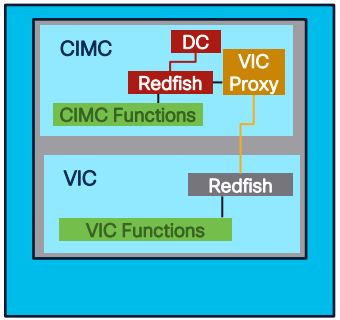

C-Series / X-Series

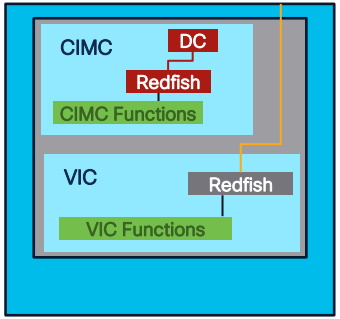

B-Series (legacy)

IMM Chassis Management Architecture

- Key Components:
  - Device Connector (“jaguar”): Enables CMC to Intersight communication. Acts as child DC to register automatically to Intersight via FI DC proxy
  - Redfish Server (IFM & IOM 2408): Provides Redfish API for managing the chassis via the Chassis Management Controller (CMC)
  - Redfish Server (IOM 220x & 2304): Provides Redfish API for managing the chassis via the Chassis Management Controller (CMC). Communicates via FI EP Proxy.

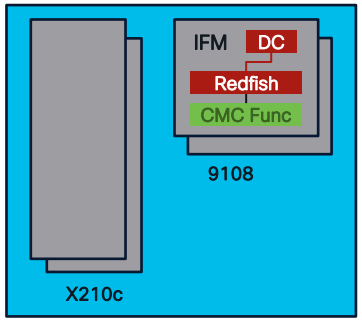

X-Series Chassis (9508)

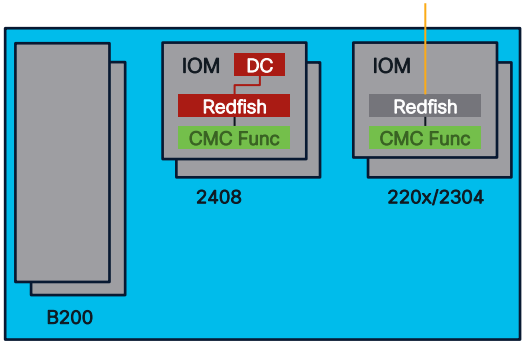

B-Series Chassis (5108)

Why Redfish?

- Standard, Scalable, Secure replacement for IPMI
- RESTful API with JSON-formatted data
  - Modern replacement for the UCS XML API
- Supports OEM/vendor extensions
  - VIC management
  - Firmware management
  - Custom actions (Reset CMOS, Clear TPM, ...)
- Available on many server platforms
  - Support for third-party servers (Dell, HPE)
- More details about Redfish on UCS: https://www.cisco.com/c/en/us/td/docs/unified_computing/ucs/c/sw/api/4_3/b-cisco-ucs-c-series-servers-rest-api-programmer-s-guide-release-4-3.html

For your reference

Device Connector versions & Upgrade

- Automatically upgraded by Intersight
  - Only requires connectivity to Intersight – no claim needed
  - DC reports its version on each startup. If version is not up to date, an immediate upgrade is executed
  - FI DC version visible:
    - under “Targets” section in Intersight UI
    - in Device Console UI (“Device Connector” tab)
    - by entering `show version` on FI SSH CLI
  - CIMC/IOM/IFM DC version visible:
    - Under “Targets” section in Intersight UI (under “Sub Target”)
  - Versioning is 1.0.11-<build datetime>.<hotfix datetime>
- Manual upgrade possible for FI DC
  - Using `update-device-connector` on FI SSH CLI
  - DC image provided by TAC
- Triggers shallow discovery upon DC restart (will be removed soon)

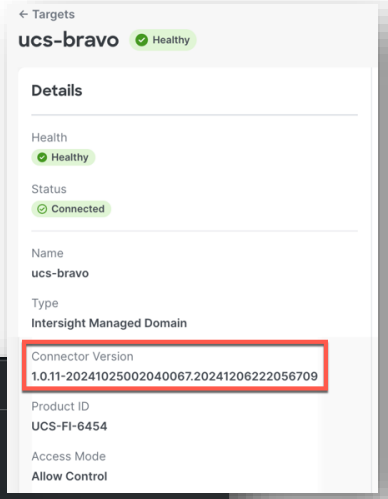

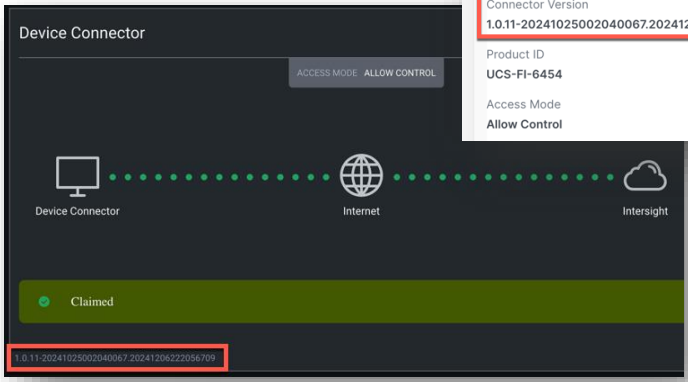

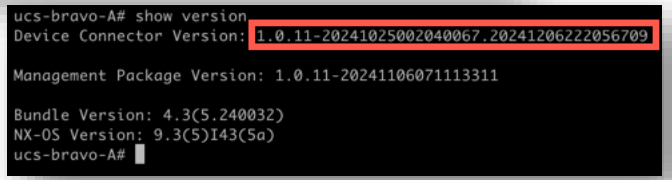

Setting up the UCS Domain

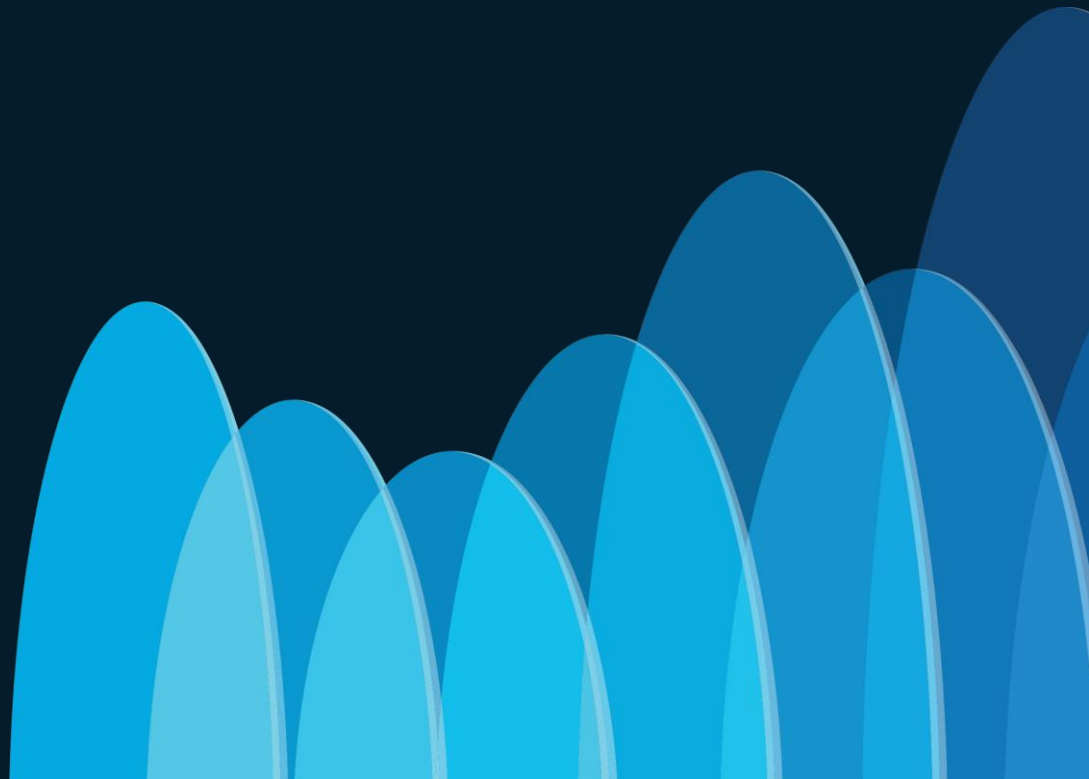

cisco Live!

Setting up a domain in Intersight Managed Mode

Management mode

```
  the Fabric interconnect and its clustering mode is performed through these steps.

  Type Ctrl-C at any time to abort configuration and reboot system.
  To back track or make modifications to already entered values,
  complete input till end of section and answer no when prompted
  to apply configuration.

first-setup:    Warning:  is EMPTY. using switch as name

   Starting GUI for initial setup.

   Switch can now be configured from GUI. Use https://172.16.105.242 and click
   on 'Express Setup' link. If you want to cancel the configuration from GUI and go back,
   press the 'ctrl+c' key and choose 'X'. Press any other key to see the installation progress from
GUI
   Note: Intersight management mode setup available through console based configuration method alone
.



^C
  Type 'reboot' to abort configuration and reboot system
  or Type 'X' to cancel GUI configuratuion and go back to console  or Press any other key to see the
 installation progress from GUI (reboot/X) ? X

  Enter the configuration method. (console/gui) ? console

  Enter the management mode. (ucsm/intersight)? intersight

  You have chosen to setup a new Fabric interconnect in "intersight" managed mode. Continue? (y/n):
y
```

The new "change-mode" CLI command

- All-in-one command to change UCS domain mode from UCSM to IMM (available from 4.3(5c))
  - Both FIs are converted to IMM simultaneously
  - FI networking configuration is kept identical (IP addresses, netmask, gateway, DNS server)
- After reboot, FIs are ready to be claimed in Intersight – no initial setup required
- All UCSM configuration is erased and is not converted to Intersight with this command
  - Leverage IMM Transition Tool if you also need to convert the configuration (profiles, templates, policies, pools)

```
UCS-BOREALIS-A# connect local-mgmt
UCS-BOREALIS-A(local-mgmt)# change-mode
All UCS configurations will be erased from both Fabric Interconnects.
Fabric Interconnects will reboot and come up as Intersight managed mode Fabric Interconnects. Are you sure? (yes/no):yes
[snip]
Removing SAM Database. Please wait....
[snip]
Removing SAM Database and rebooting peer switch.
[snip]
Database is cleaned up. Rebooting....
[snip]
Starting Intersight managed UCS Processes..
Cisco UCS 6500 Series Fabric Interconnect
UCS-BOREALIS-A login:
```

After Initial Setup…

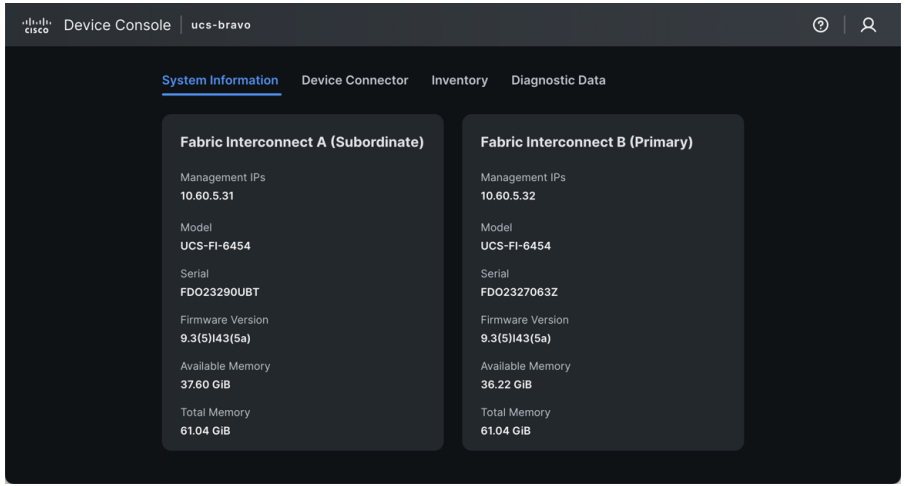

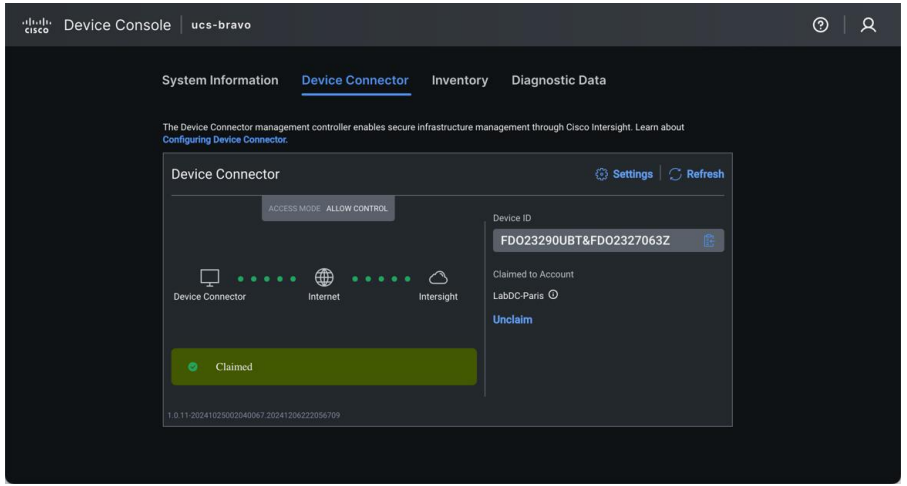

Device Console (part of "`dove`" Management Package) available from both FI IPs:

- Configuration of "`jaguar`" Device Connector (DNS, NTP, Proxy) (automatically synced between FI peers)
- Copy Device ID & Claim Code for SaaS claiming (Appliance claiming uses Device Console admin credentials)
- Inventory with Server Power actions, Locator LED & vKVM Launch (empty at first) + Redfish-based API Explorer
- Tech-Support bundle collection of connected devices

IMM Domain cluster HA

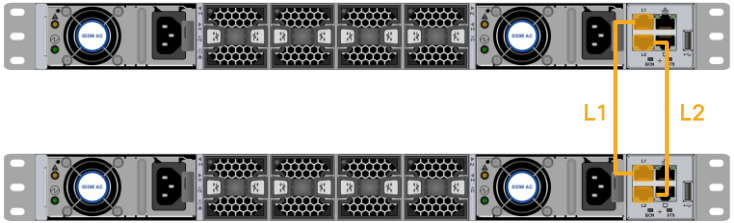

- Only cluster mode supported in IMM (no standalone FI)
- FIs are clustered using L1-L2 ports for:
  - Synchronizing DC configuration (Proxy, DNS, NTP, etc.)
  - Synchronizing firmware images during upgrades
  - Synchronizing MAC tables between FIs (used for Fabric Failover of VM MACs)
- Primary / Subordinate role only used for managing the DC configuration (only Primary can write to the DC database)
  - No impact to server management or FI firmware upgrade order
  - Not shown in Intersight UI as it is not “important”
- Both FI DCs have an active connection to Intersight
  - Configuration specific to each FI sent through their respective DC
  - No NX-OS switch config sync between FIs

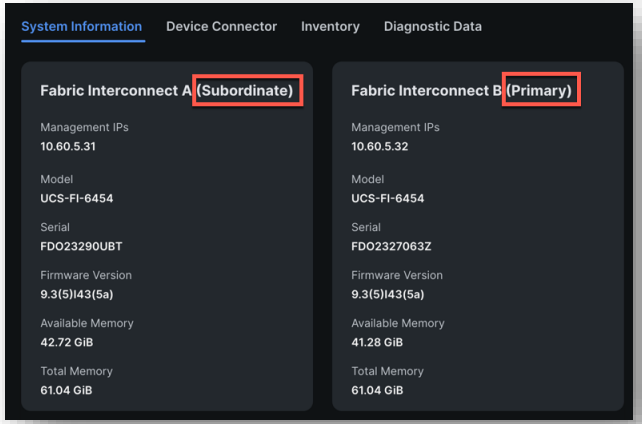

Claiming the IMM Domain

- DC uses “CloudDnsList” and “CloudDns” (fallback) values to determine which Intersight instance to connect to
  - Defaults to “svc.intersight.com”
- With Intersight SaaS
  - Using Device ID and Claim Code from DC
- With Intersight Appliance
  - Using IMM FI admin account credentials
  - Automatically sets the “CloudDns” value to the appliance FQDN
  - Automatically uses the Device ID and Claim Code
- FIs (`network.Element` MOs) automatically inventoried upon successful claim

DC *connector.db*

```
{
   "CloudDns": "svc.intersight.com",
   "CloudDnsList": [
      "svc.intersight.com",
      "svc-static1.intersight.com",
      "svc.ucs-connect.com",
      "svc-static1.ucs-connect.com"
   ],
   "ConnectionId": "ea02202b485c70a9d6fa",
   "ConnectionState": "Connected",
   "DnsLatency": "14.877141ms",
   "ResolvedAddresses": [
      "64.103.36.133"
   ],
   "ConnectedAddress": "64.103.36.133:80",
   "ClientIPAddress": "10.60.5.31:43977",
   "ProxyType": "ProxyConfigured",
   [snip]
}
```

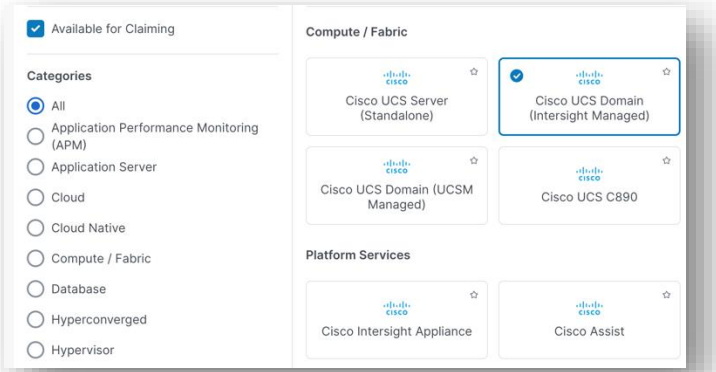

FI Device Connector REST API

- Login to **https://<fi_ip>/Login**
  - Send POST with credentials in raw Body
  - Returns SessionId value for subsequent API calls

```
{
   "User": "admin",
   "Password": "!Cisco123"
}
```

```
{
   "Status": 200,
   "Description": "",
   "SessionId": "YWQDjMTGYIB4Dmd1EVi-9aBrdOfgeh8y587axC42HGE%3D",
   "User": "admin",
   "Role": "Account Administrator"
}
```

- Send GET to **https://<fi_ip>/connector/<Endpoint>**
  - **DomainName**: IMM Domain Name
  - **SystemInfo**: FI System info (content visible in Device Console)
  - **DeviceConnections**: Device Connector connection to Intersight + Proxy config
  - **DeviceIdentifiers**: Device ID (for claiming)
  - **SecurityTokens**: Claim Code – only available if connectivity to Intersight is established
- Reset Device Connector config to factory defaults
  - Send PUT to **https://<fi_ip>/connector/DeviceConnections** with raw Body

```
{
    "CloudDns": "svc.intersight.com",
    "ForceResetIdentity": true
}
```

UCS Domain Profile policy structure

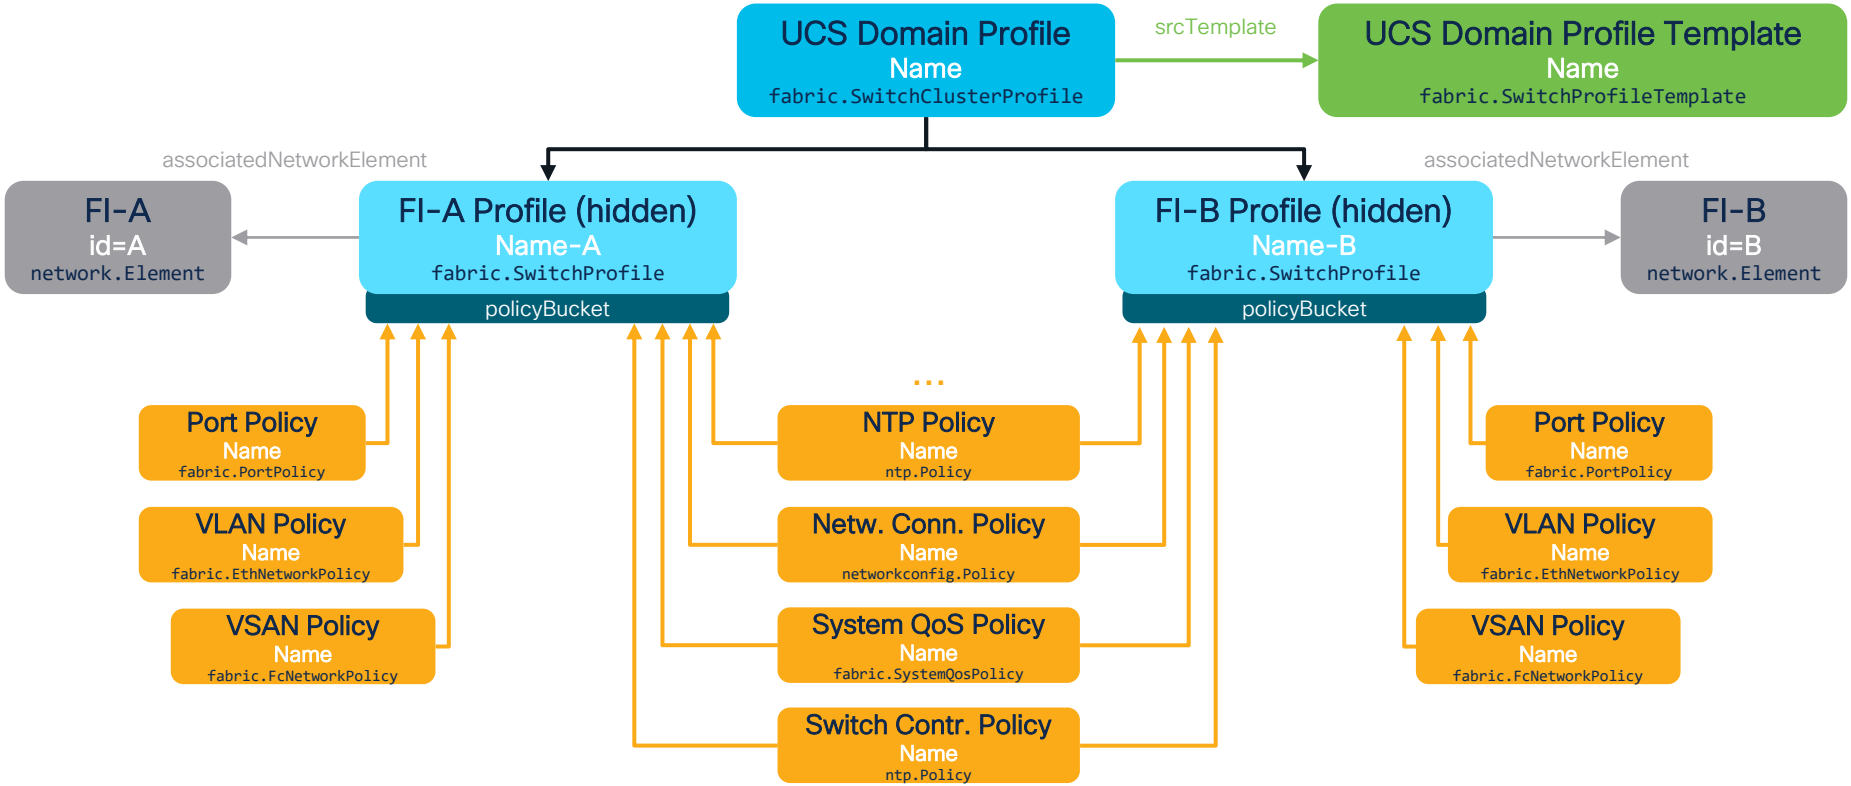

Deploying the Domain Profile

- Always initiated by user
  - Never automatic based on policy change
- Initiates 2 workflows (one for each Fabric Interconnect)
  - Validates each policy prior to deployment
- Configuration attributes requiring FIs reboot:
  - Changing Unified Ports range
  - Changing Ethernet or FC switching mode (End-Host/Switch)
- Reboot always sequential, starting with FI-B then FI-A (same as for firmware upgrades)
  - No user acknowledgment – allows for automated initial deployment
  - Further changes (not very common) should be done one fabric at a time
- Unassign of Domain Profile is non-disruptive
  - Warning: Unclaim of IMM Domain Target cleans up the configuration from the FIs which is disruptive to the data paths

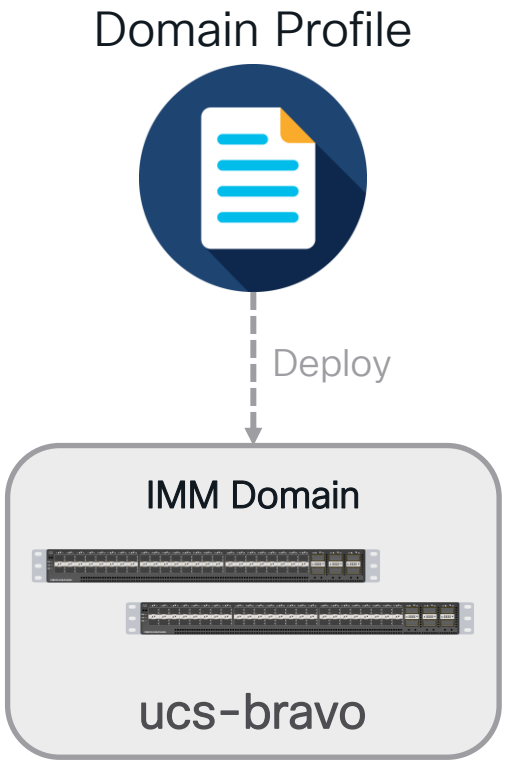

Fabric Interconnect NX-OS Config Replay

- FI configuration needs to persist state across reboots
  - Allows system to recover autonomously after power outage (without Intersight)
- Config is persisted in the FI-DC database, not in NX-OS
  - On every FI boot, the NX-OS configuration is empty
- Once FI-DC is operational, it replays the device model in proper order to reconfigure NX-OS
- During infra firmware upgrades, after FI reboot:
  - Once FI-DC is operational, FI is operational and reachable
  - However, vNIC/vHBA data paths take time to restore as Config Replay needs time to complete (uplink ports are configured last to reduce FI load)

IMM Domain (Device Console) SSL certificates

- Self-signed certificate (default)
  - Used by **dove** for Web UI
  - Auto-generated for 1 year with 2,048 bits key
  - Can be re-generated by using the **generate-self-signed-certificate** CLI command
- Custom certificate in roadmap

```
ucs-bravo-A# generate-self-signed-certificate
hostname is ucs-bravo-A
Successfully generated the self-signed-certificates
Successfully restarted the web-server
ucs-bravo-A#
```

- Add Trusted Root CA to IMM domain
  - Useful for connecting to Secure LDAP server
  - Attach Certificate Management Policy to UCS Domain Profile

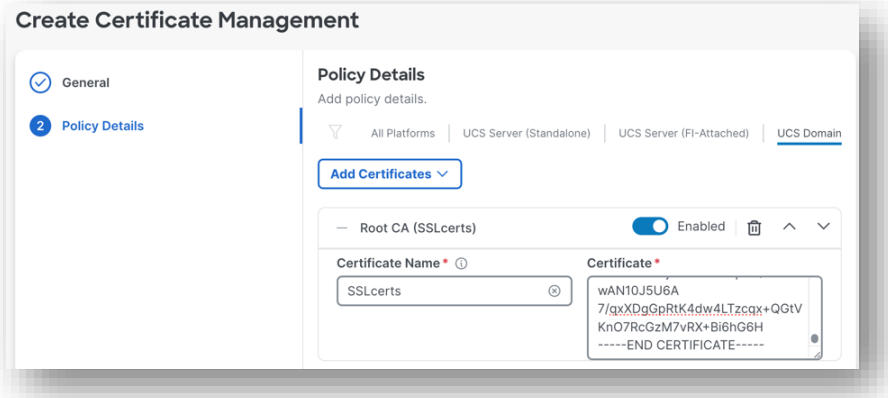

Discovering Hardware

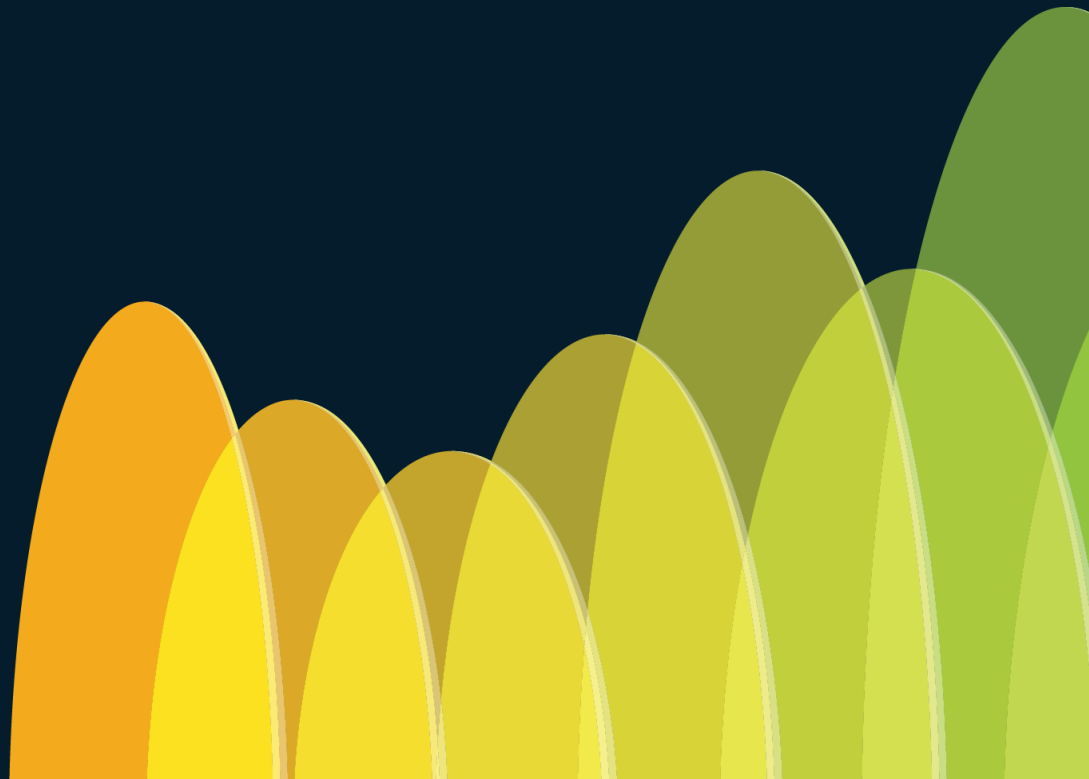

cisco Live!

Chassis & Blades

UCS Chassis Internal Connectivity (1)

IFM 25G

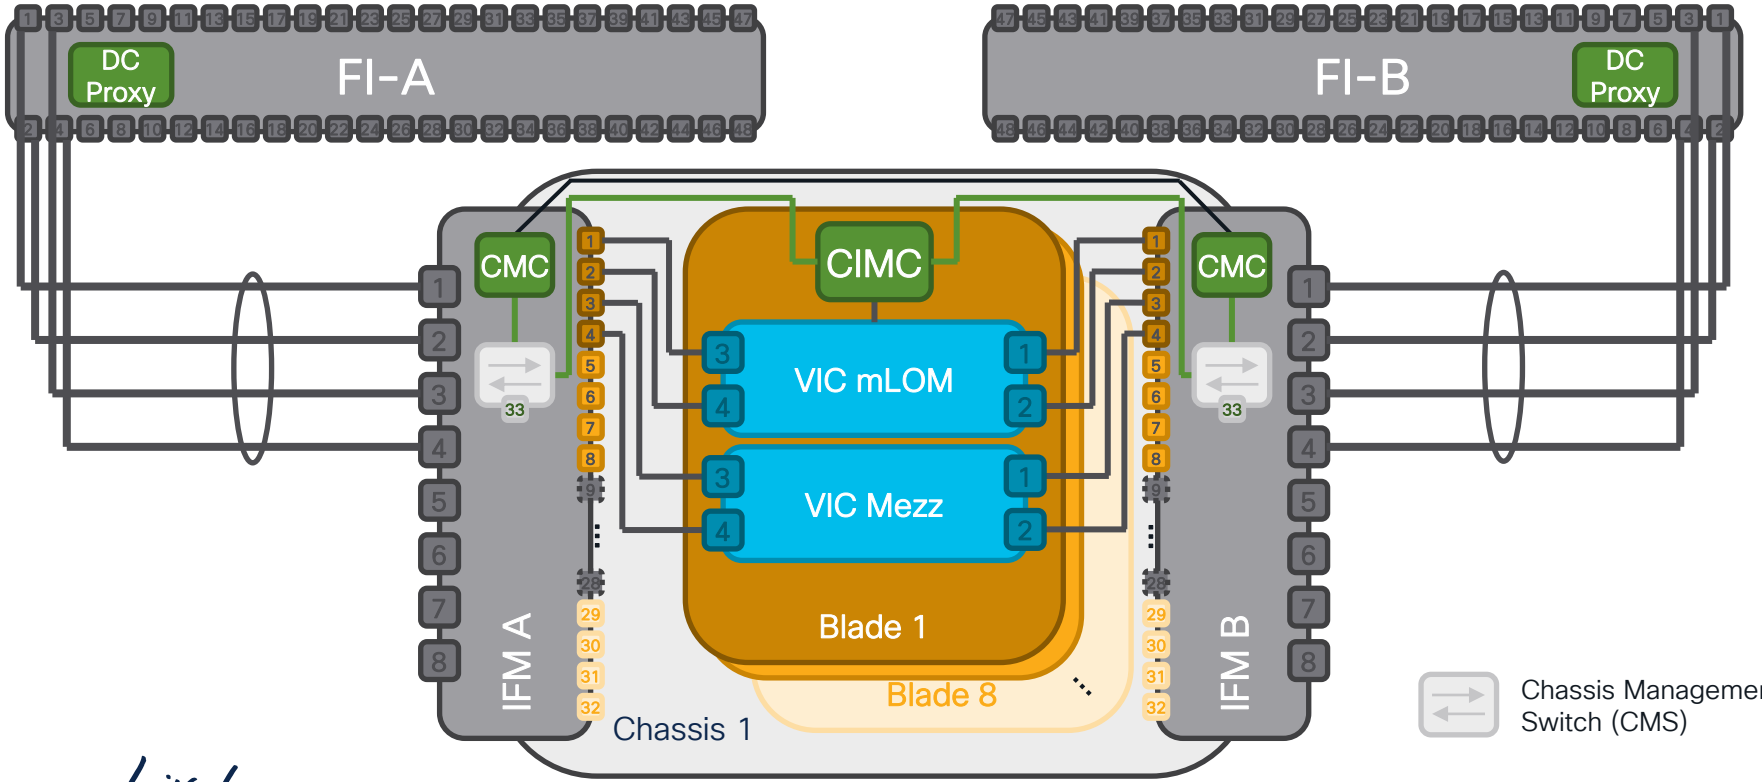

UCS Chassis Internal Connectivity (2)

IFM 25G

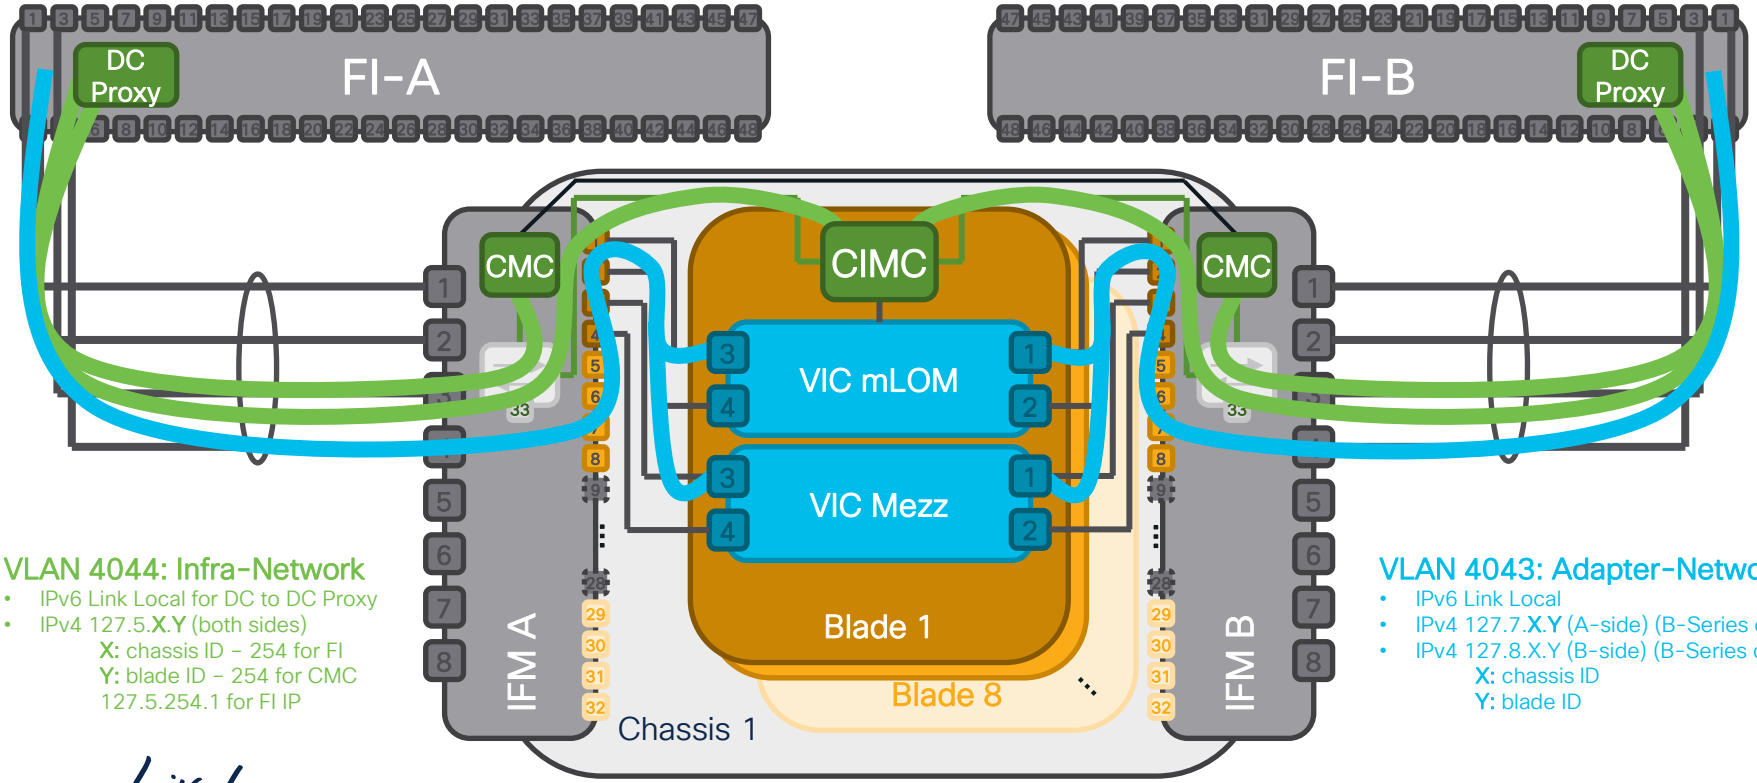

VLAN 4044: Infra-Network

- IPv6 Link Local for DC to DC Proxy
- IPv4 127.5.**X.Y** (both sides)
  - **X:** chassis ID – 254 for FI
  - **Y:** blade ID – 254 for CMC
  - 127.5.254.1 for FI IP

VLAN 4043: Adapter-Network

- IPv6 Link Local
- IPv4 127.7.**X.Y** (A-side) (B-Series only)
- IPv4 127.8.X.Y (B-side) (B-Series only)
  - **X:** chassis ID
  - **Y:** blade ID

Chassis Discovery (1)

1. Once port has Server role, NX-OS config is initially set to:

```
interface Ethernet1/1
  description Server
  no pinning server sticky
  switchport mode trunk
  switchport trunk allowed vlan 4044
  no shutdown
```

Allows chassis or rack server discovery

2. Uses Satellite Discovery Protocol (SDP) to detect IOM/IFM presence
   - Includes Model, Serial, Vendor info
3. IOM Presence event generated by NX-OS is sent to Intersight (`gobi`) via INA
   - Creates `equipment.ChassisIdentity` object to manage chassis lifecycle
4. Chassis Discovery workflow is triggered for each server port
   - Creates `equipment.Chassis` & `equipment.IoCard` objects

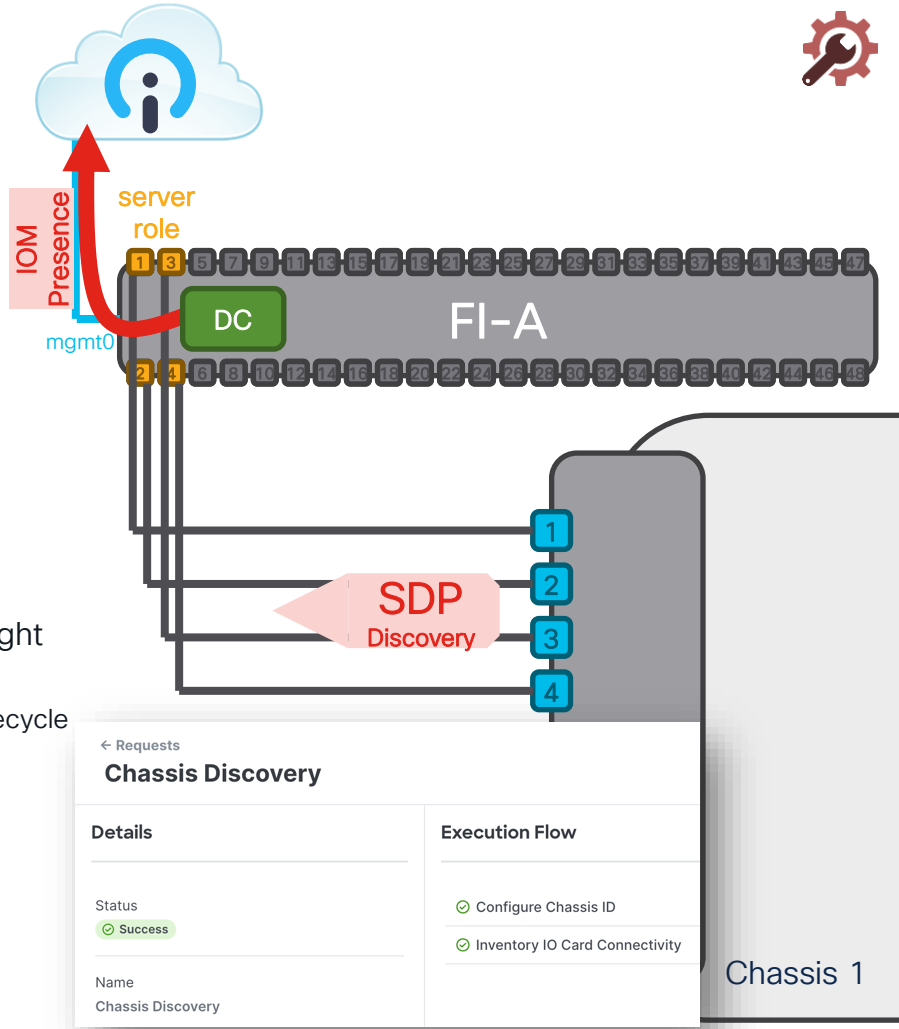

Chassis Discovery (2)

5. Intersight configures fabric port-channel on FI through INA request:

```
interface Ethernet1/1
  description Server
  no pinning server sticky
  switchport mode fex-fabric
  priority-flow-control mode on
  fex associate 1
  fec rs-fec
  channel-group 1025
  no shutdown
```

Automatic Port-Channel numbering

```
interface port-channel1025
  description FabricPc
  switchport mode fex-fabric
  no pinning server sticky
  priority-flow-control mode on
  fex associate 1 chassis-serial FOX2213P2VY
     module-serial FCH242070PN module-slot left
```

Chassis S/N

IOM S/N

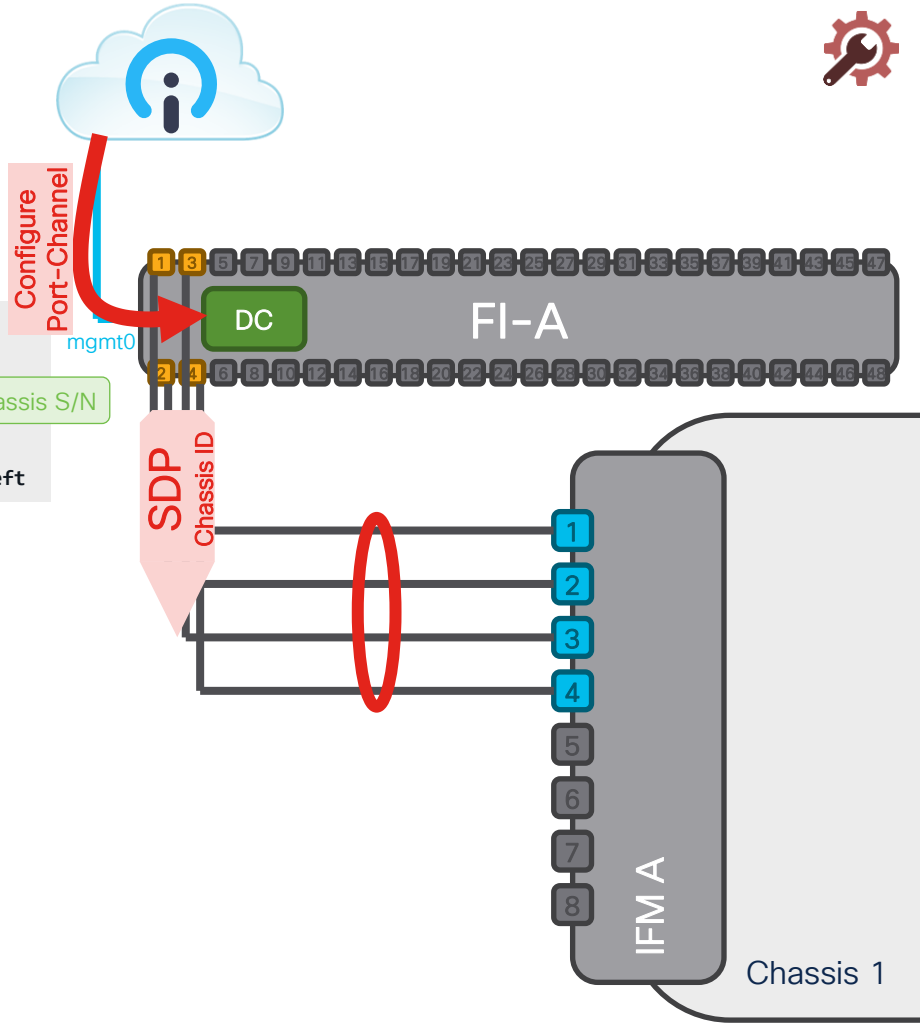

- Chassis/FEX ID gets assigned to IOM/IFM via SDP, along with FI IP & Management Mode (UCSM / IMM) information
- FI NX-OS checks IOM/IFM software compatibility and triggers automatic firmware upgrade/downgrade if required
- Use `show fex <id> detail` under NX-OS to track the discovery progress & IOM/IFM auto-upgrade

Chassis Discovery (3)

6. IOM Online event generated by NX-OS is sent to Intersight (`gobi`) via INA

7. Chassis Inventory workflow is triggered for each IOM/IFM
   - Configures internal port for communication with CMC:

```
interface Ethernet1/1/33
  no pinning server sticky
  switchport mode trunk
  switchport trunk native vlan 4044
  switchport trunk allowed vlan 4044
  no shutdown
```

   - CMC DC starts auto-claim (registration) to Intersight via DC Proxy
   - Intersight (**`dejavu`**) inventories chassis (FANs, PSUs, IOMs and Blade slots) through Redfish API

8. `compute.BladeIdentity` object is created for each discovered blade based on presence pin:

```
fex-1# show platform software cmcctrl blade inventory
Blade,Status,BladeClass,BmcSlot,Occupied,Make,Product,SerialNo
B1,DISCOVERED,0,B1,occupied,"Cisco Systems Inc","UCSX-210C-M6","FCH271278ZV"
B2,DISCOVERED,0,B2,occupied,"Cisco Systems Inc","UCSX-210C-M6","FCH271278Z5"
B3,DISCOVERED,0,B3,occupied,"Cisco Systems Inc","UCSX-210C-M6","FCH2712795F"
B4,DISCOVERED,0,B3,occupied,"Cisco Systems Inc","UCSX-440P","FCH261171W5"
B5,DISCOVERED,0,B5,occupied,"Cisco Systems Inc","UCSX-210C-M7","FCH263873H9"
B6,DISCOVERED,0,B6,occupied,"Cisco Systems Inc","UCSX-210C-M7","FCH264272G6"
B7,UNKNOWN,0,B7,empty,"","",""
B8,UNKNOWN,0,B8,empty,"","",""
```

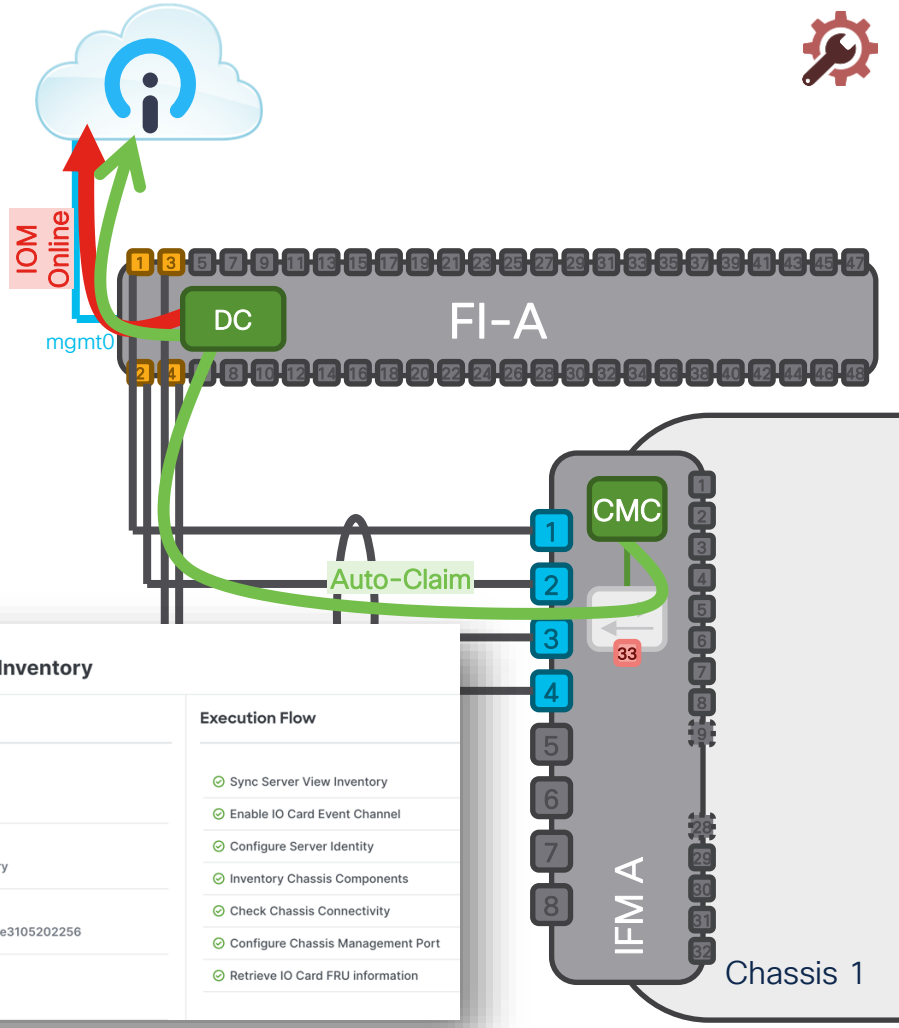

Chassis Lifecycle Operations

- **Chassis Decommission:** (disruptive to all servers in the chassis)

  Deletes chassis inventory MOs and brings IOMs/IFMs offline by moving server ports from **`fex-fabric`** mode to **`trunk`** mode. `equipment.ChassisIdentity` MO is is preserved and marked as decommissioned.

- **Chassis Remove:**

  Removes chassis from inventory. Can only be performed if there are no server ports connected to the chassis. Deletes `equipment.ChassisIdentity` MO.

- **Chassis Rediscover:** (non-disruptive to all servers)
  - Rebuilds connectivity between FI and IOMs/IFMs
  - Triggers Chassis Discovery (IOM Presence) workflow for each port & cleans up ports no longer connected to chassis
  - Triggers Chassis Inventory (IOM Online) workflow to re-inventory all chassis components

- **Chassis Recommission:**

  Recommissions a decommissioned chassis by triggering Chassis Rediscover workflow. Can be used to change chassis ID.

Chassis Troubleshooting Operations

- **IOM/IFM Reset:**

  Allows an individual IOM/IFM to be reset.

- **Peer IOM/IFM Reset:**

  Allows resetting an IOM/IFM from its peer. Can be used to recover a non-responding IOM/IFM instead of doing a physical IOM/IFM reseat. Leverages internal link between IOMs/IFMs in the chassis to trigger the reset.

- **Power Cycle Chassis Slot:**

  Allows power cycling a server slot within the chassis. Can be used to recover a non-responding CIMC instead of doing a physical server reseat.
  Disruption Warning: Brings down power to both the CIMC and the server simultaneously.

Blade Discovery

1. Blade Discovery workflow is triggered for each blade based on presence pin
2. Blades Device Connectors start auto-claim (registration) to Intersight via DC Proxy
   - CIMC DC (`jaguar`) inherits the FI DC configuration (CloudDns, proxy) and auto-claims to the same Intersight account/instance
   - Uses Link-Local IPv6 address for CIMC to FI communications
3. IOM/IFM HIF (server) ports get configured through INA request:
   - Uses automatic port-channels for agregating multiple VIC ports to the same IOM/IFM

```
interface Ethernet1/1/1
  description BladeServer
  no pinning server sticky
  switchport mode vntag
  switchport trunk allowed vlan 4043
  channel-group 1290
  no shutdown
```

Enable VNTag support to VIC adapter

```
interface port-channel1290
  description HostPC
  switchport mode vntag
  switchport trunk allowed vlan 4043
  no pinning server sticky
```

4. Retrieves inventory of all components (deep discovery)
   - Powers on server and waits for BIOS POST to complete
   - Uses Redfish API calls to blade CIMC DC (`jaguar`)

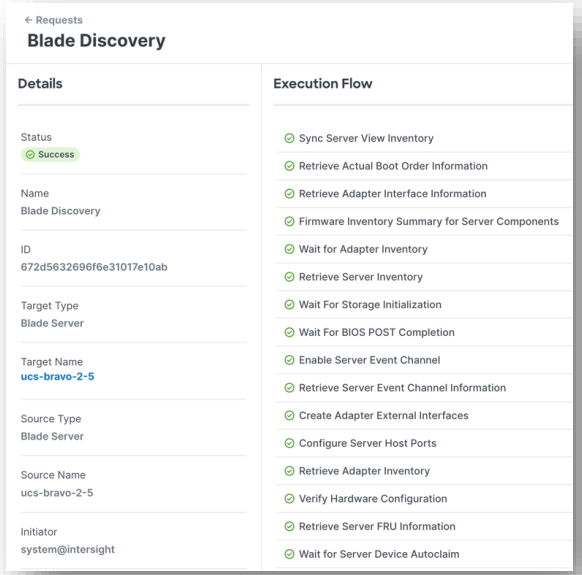

“Deep” vs “Shallow” Discovery

- Attribute of the Discovery operation (workflow) to skip some tasks depending on the context.
- “Deep” discovery: full inventory of the server including SMBIOS information & VIC’s interfaces. Always powers on host at the beginning of the workflow (if not already powered on).
  - Triggered by Initial discovery & user-initiated actions (including server “Rediscover”)
  - Automatically upgrades VIC firmware if version does not support IMM (RMA case)
- “Shallow” discovery: hardware-based inventory of the server to detect any hardware change on the server. Does NOT change host power state at any point.
  - Triggered after CIMC reboot or CIMC DC restart (will be removed in the future)
  - Triggered after Chassis Inventory workflow for already discovered blades
- Non-disruptive
  - Blade/Rack Server Discovery is always performed “out of band” – i.e. without requiring a server reboot or starting a specific OS (unlike PnuOS in UCSM)

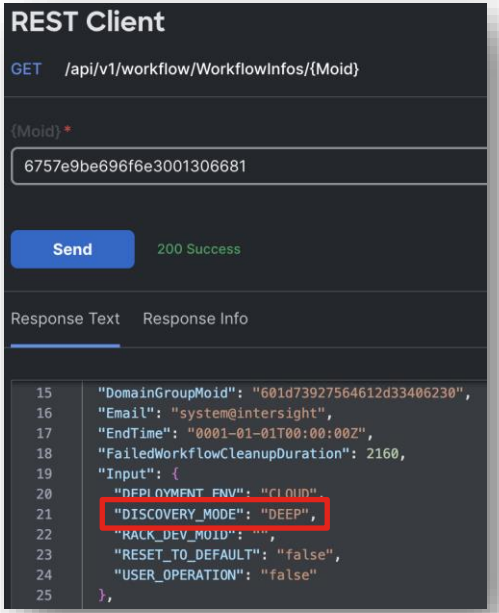

Blade Lifecycle Operations

- **Blade Decommission:** (disruptive to the server)

  Deletes blade inventory MOs, device registration and unconfigures blade-specific FI configuration (HIF ports). `compute.BladeIdentity` MO is is preserved and marked as decommissioned. Resets server to factory defaults.

- **Blade Remove:**

  Removes blade from inventory. Can only be performed if blade is in decommissioned state. Deletes `compute.BladeIdentity` MO.

- **Blade Rediscover:** (non-disruptive to the server)

  Triggers Blade Discovery (“deep”) workflow to re-inventory Blade components

- **Blade Recommission:**

  Recommissions a decommissioned blade by triggering Blade Discovery workflow.

- **Blade Secure Erase:** (disruptive to the server and its data)

  Deletes all data on the BIOS, CIMC, NVRAM, DIMM, eMMC, VIC and local drives. Can take several hours. Complies with EU Lot 9 and NIST SP 800-88 standards for data sanitization.

Blade Troubleshooting Operations

- **Reboot Management Controller:**

  Allows a CIMC (and corresponding DC) to be reset. Automatically triggers “shallow” discovery after CIMC DC is back online.

- **Reset CMOS:**

  Resets the CMOS memory that stores BIOS settings. Useful to troubleshoot boot or hardware compatibility issues. Only available when server is powered off. For reset to complete, server must be powered on after the operation.

- **Reset vKVM:**

  Triggers a workflow that restarts the vKVM service running on the CIMC. Useful to troubleshoot when launch of vKVM is failing or vKVM becomes unresponsive. Supported with firmware version 4.2(1a) and above.

- **Reset Memory Errors:**

  Resets all memory error counters of all DIMMs to 0. All disabled DIMMs due to prior uncorrectable errors (UECC) will be re-enabled on next boot if they don’t encounter further errors during POST.

- **Clear System Event Logs (SEL):**

  Clears all SEL Logs from the CIMC memory. Useful when they are full. Contain records of hardware events and failures, and other system critical events.

X-Series Chassis Power Allocation for Initial Discovery

For your reference

- Default Redundancy Policy for X-Series chassis: N+1
  - Changed from default **Grid** since Nov 2024 in SaaS (Appliance 1.1.2-0)
  - Allows for up to 16,100W power budget during discovery (including Power Extended Mode)
- Initial Power budget allocation (for initial discovery only):
  - Unused slot: 25W (to power on CIMC upon server insertion)
  - X210c M6: 1300W (1480W with X440p - depends on config)
  - X210c M7: 1610W
  - X410c M7: 3175W
  - Chassis IFMs, XFMs, Fans: 2084W
- After initial discovery, allocates profiled max or catalog max
- X-Series Power & Cooling detailed session: BRKCOM-2933 https://www.ciscolive.com/c/dam/r/ciscolive/global-event/docs/2024/pdf/BRKCOM-2933.pdf

```
fex-1# show platform software cmcctrl power all
[snip]
Chassis power allocation report:
-------------------------------------------------------------------------------
Chassis powerlimit: 16100                                                      |
-------------------------------------------------------------------------------
Chassis blade available budget: 15615                                          |
-------------------------------------------------------------------------------
Chassis blade allocated budget: 14089                                          |
-------------------------------------------------------------------------------
ebrake_min_req_psu: 0 | prereq_ebrake =             N/A                       |
-------------------------------------------------------------------------------
        |presence | init_pwr| dyn_pwr | ok_to_pwr_on | ok_to_pwr_prof | fault   |
-------------------------------------------------------------------------------
Blade1  | Present |   1300W |   1300W | Yes          | No             | No      |
-------------------------------------------------------------------------------
Blade2  | Present |   1300W |   1300W | Yes          | No             | No      |
-------------------------------------------------------------------------------
Blade3  | Present |   1480W |   1480W | Yes          | No             | No      |
-------------------------------------------------------------------------------
Blade4  | ---     |         |         |              |                |         |
-------------------------------------------------------------------------------
Blade5  | Present |   1610W |   1610W | Yes          | No             | No      |
-------------------------------------------------------------------------------
Blade6  | Present |   1610W |   1610W | Yes          | No             | No      |
-------------------------------------------------------------------------------
Blade7  | Present |   3175W |   3175W | Yes          | No             | No      |
-------------------------------------------------------------------------------
Blade8  | ---     |         |         |              |                |         |
-------------------------------------------------------------------------------
SubTotal|         |  10475W |  10475W |
-----------------------------------------
Ch Base |         |   2084W |    485W |
-----------------------------------------
Total   |         |  12559W |  10960W |
-----------------------------------------
```

Rack Servers

UCS Rack Server Internal Connectivity (1)

Direct Connect (without FEX) with VIC 4x25G

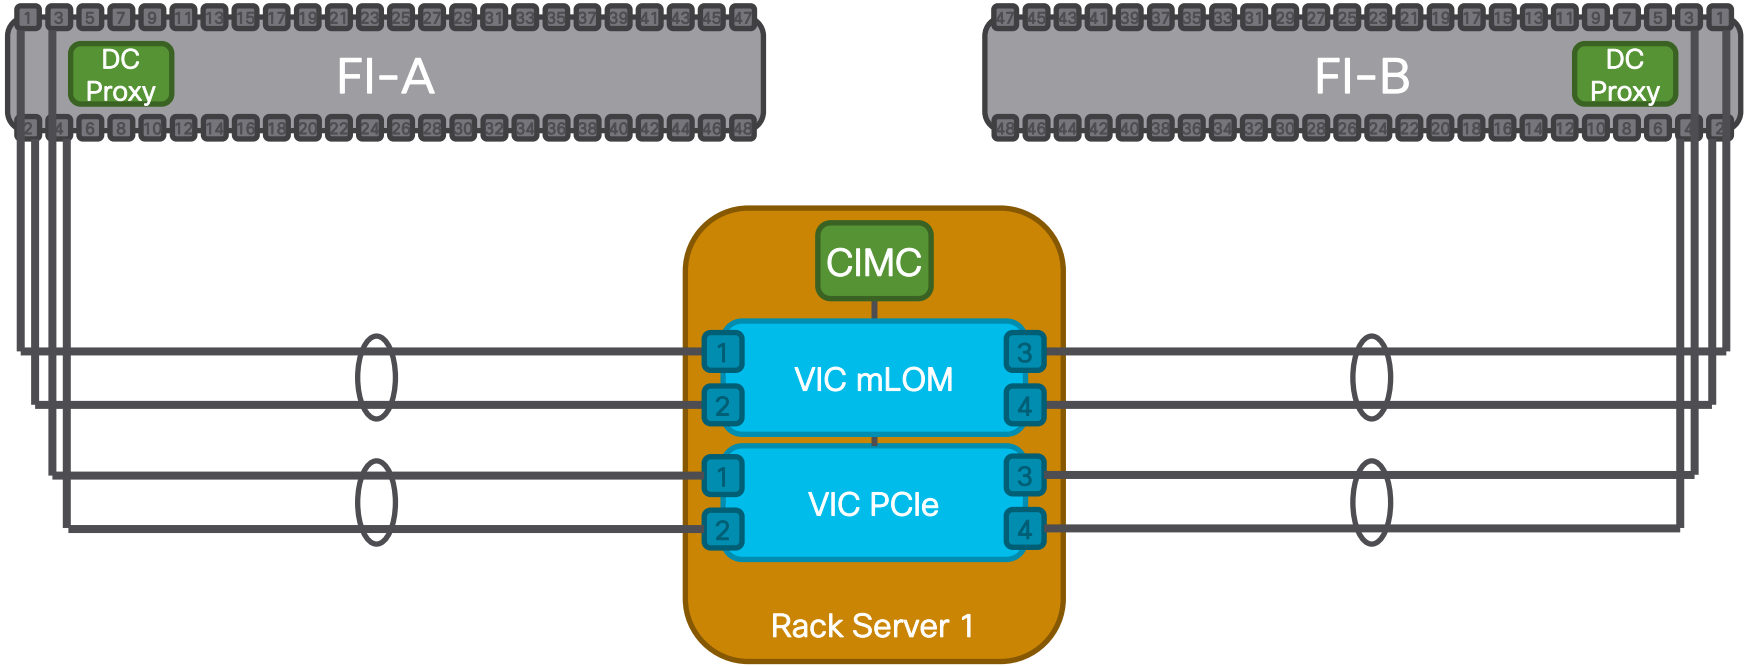

UCS Rack Server Internal Connectivity (2)

Direct Connect (without FEX) with VIC 4x25G

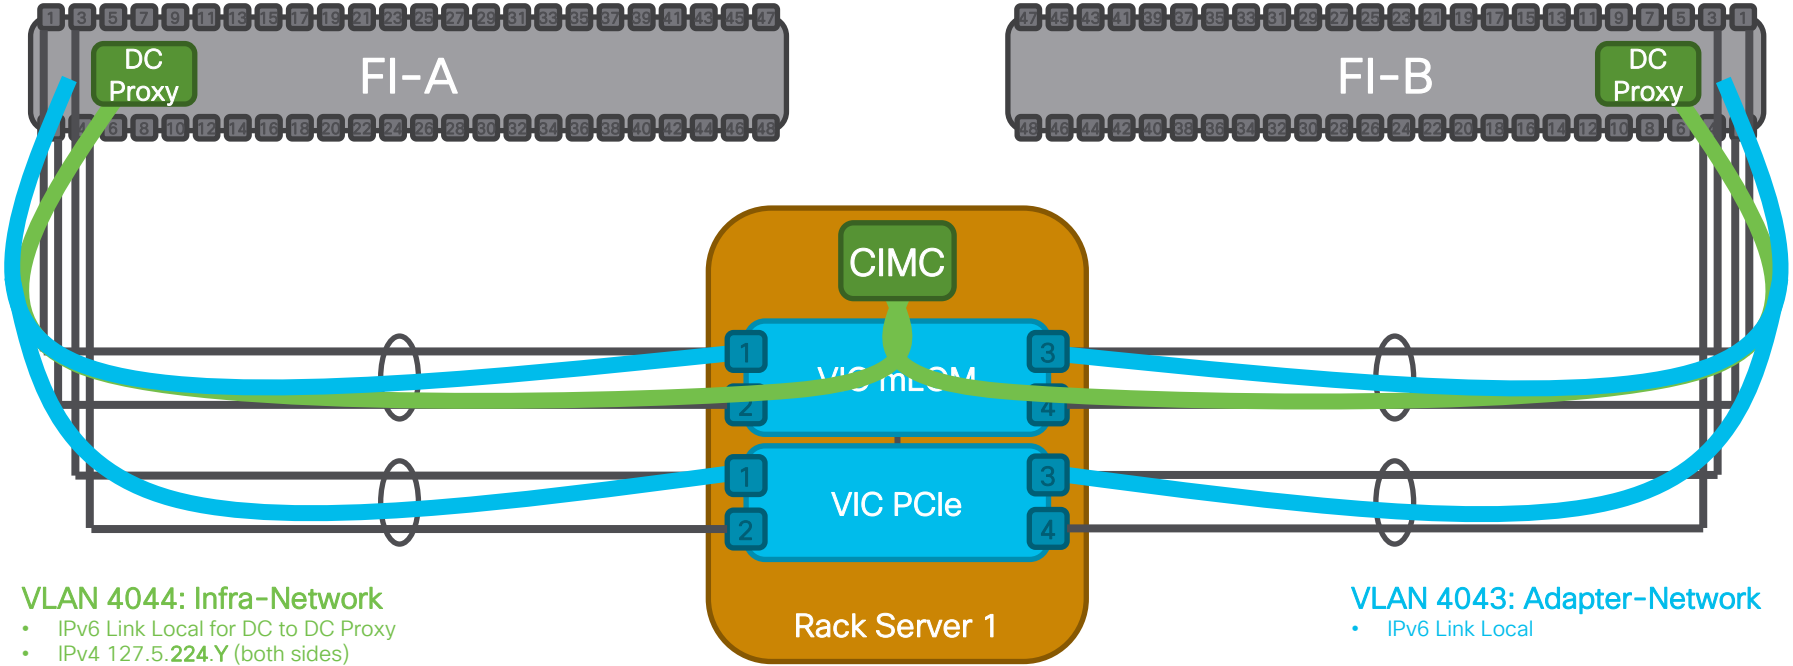

VLAN 4044: Infra-Network

- IPv6 Link Local for DC to DC Proxy
- IPv4 127.5.**224**.**Y** (both sides)
  Y: rack server ID
  127.5.254.1 for FI IP

VLAN 4043: Adapter-Network

- IPv6 Link Local

Rack Server Discovery (1)

1. Once port has Server role, NX-OS config is initially set to:

```
interface Ethernet1/1
  description Server
  no pinning server sticky
  switchport mode trunk
  switchport trunk allowed vlan 4044
  no shutdown
```

Allows chassis or rack server discovery

2. mLOM VIC uses Link-Layer Discovery Protocol (LLDP) to detect rack server presence
   - Includes Model, Serial, Vendor info of server + VIC
3. Rack Presence event generated by NX-OS is sent to Intersight (`gobi`) via INA
   - Creates `compute.RackUnitIdentity` object to manage rack lifecycle
   - Rack ID gets assigned (unique to the IMM domain)

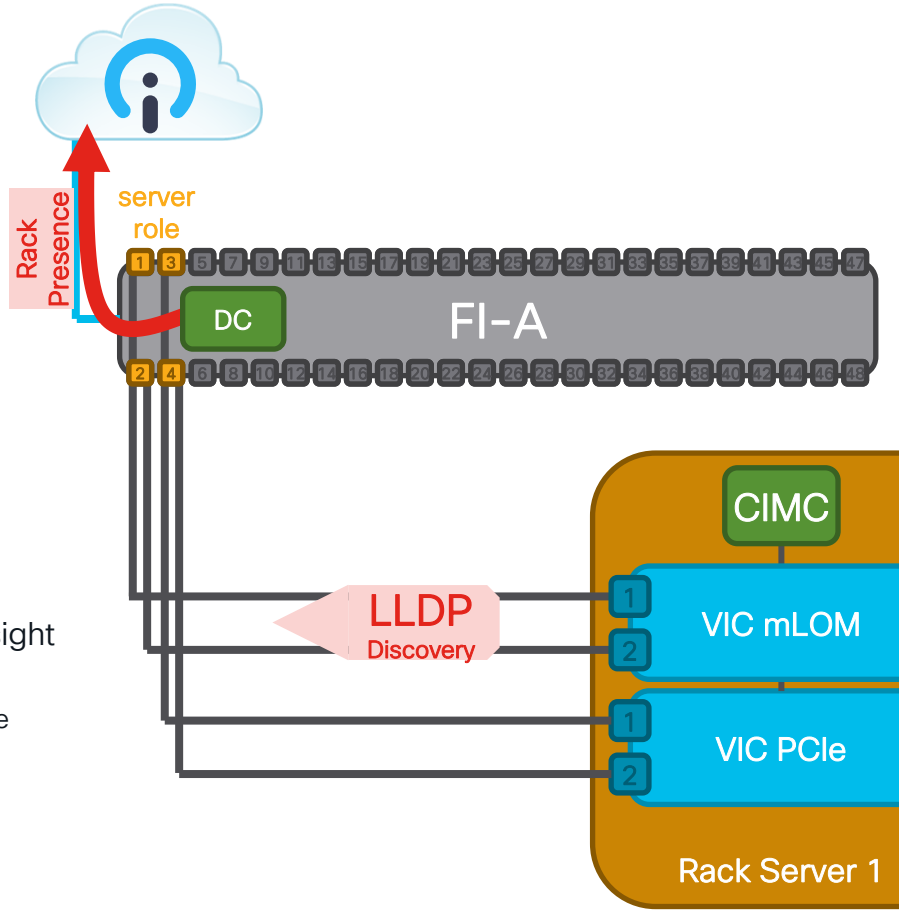

Rack Server Discovery (2)

4. Intersight configures VNTag port-channel on FI through INA request:

```
interface Ethernet1/1
  description Server
  chassis associate 1 chassis-serial WZP23030BZX
    module-serial FCH22337SHH port-profile dummy module-side left
  no pinning server sticky
  switchport mode vntag
  switchport trunk allowed vlan 4043
  channel-group 1280
  no shutdown
```

Automatic Port-Channel numbering

```
interface port-channel1280
  description Server
  switchport mode vntag
  chassis associate 1 chassis-serial WZP23030BZX
    module-serial FCH22337SHH port-profile dummy module-side left
  no pinning server sticky
  switchport trunk allowed vlan 4043
```

Server S/N

VIC S/N

5. FI uses LLDP TLV to send info to server (rack ID, vntag mode, etc.)

6. CIMC changes mode to IMM and enables VNTag on the VIC adapter(s)

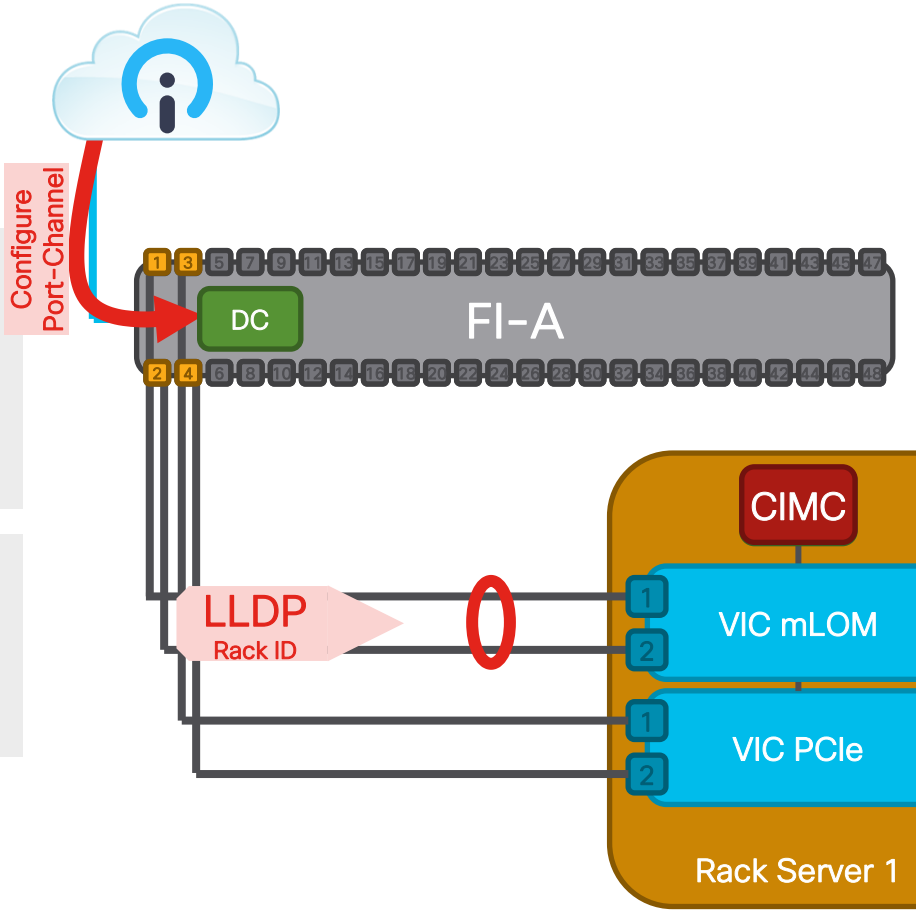

Rack Server Discovery (3)

7. FI uses Veth interface to communicate with CIMC via VIC:

```
interface Vethernet32768
  description veth-32768
  no pinning server sticky
  switchport mode trunk
  switchport trunk allowed vlan 4044
  no hardware vethernet mac filtering per-vlan
  bind interface port-channel1280 channel 65535
  no shutdown
```

Static VNTag channel ID

8. Device Connector starts auto-claim (registration) to Intersight via DC Proxy
   - CIMC DC (**jaguar**) inherits the FI DC configuration (CloudDns, proxy) and auto-claims to the same Intersight account/instance
   - Uses Link-Local IPv6 address for CIMC to FI communications

9. Rack Server Discovery workflow is triggered
   - Powers on server and waits for BIOS POST to complete
   - Retrieves inventory of all components (deep discovery)
   - Uses Redfish API calls to rack server CIMC DC (**jaguar**)
   - Configures PCIe VIC adapter connectivity (using LLDP)

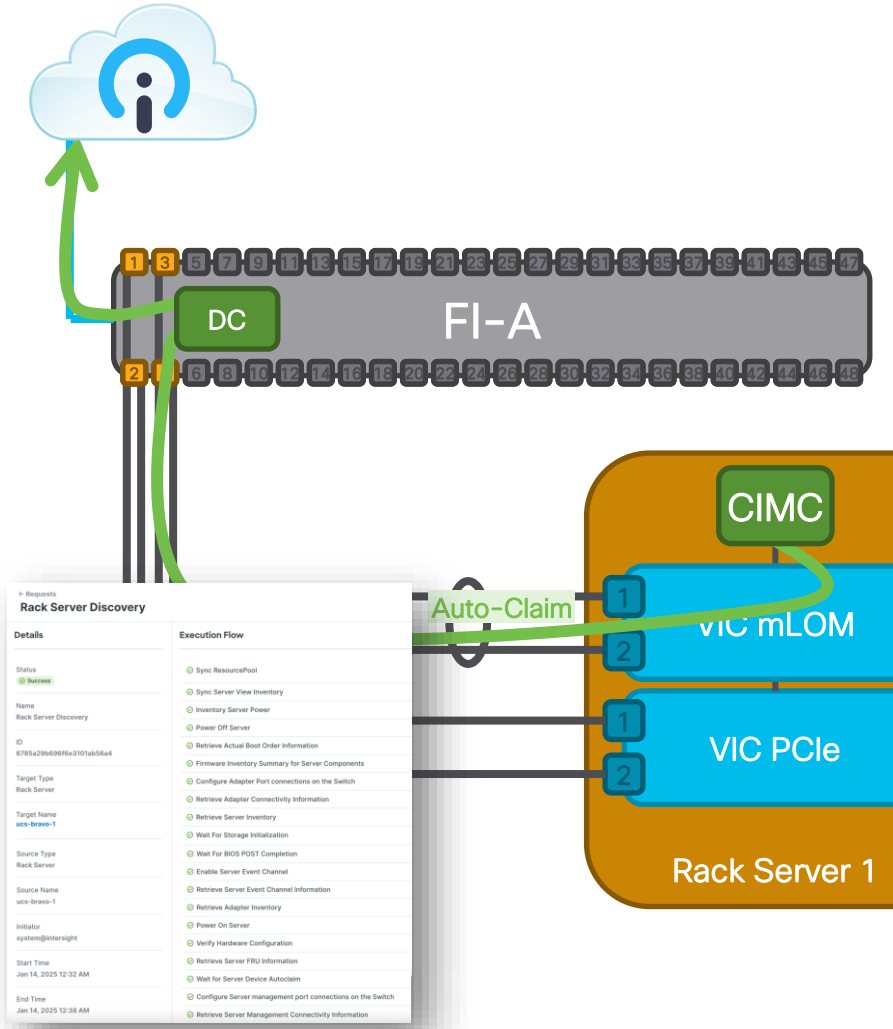

Rack Server Lifecycle Operations

- **Rack Decommission:** (disruptive to the server)

  Deletes rack inventory MOs, device registration and unconfigures rack-specific FI configuration by moving server ports from **`vntag`** mode to **`trunk`** mode. Also restores the CIMC to factory defaults. `compute.RackUnitIdentity` MO is is preserved and marked as decommissioned.

- **Rack Remove:**

  Removes rack from inventory. Can only be performed if rack is in decommissioned state. Deletes `compute.RackUnitIdentity` MO.

- **Rack Rediscover:** (non-disruptive to the server)
  - Triggers Rack Discovery workflow (“deep”) for each port
  - Rebuilds connectivity between FI and rack server & cleans up ports no longer connected to rack server

- **Rack Recommission:**

  Recommissions a decommissioned rack by triggering Rack Discovery workflow.

Rack Server Troubleshooting Operations

- **Reboot Management Controller:**

  Allows a CIMC (and corresponding DC) to be reset. Automatically triggers “shallow” discovery after CIMC DC is back online.

- **Reset CMOS:**

  Resets the CMOS memory that stores BIOS settings. Useful to troubleshoot boot or hardware compatibility issues. Only available when server is powered off. For reset to complete, server must be powered on after the operation.

- **Reset vKVM:**

  Triggers a workflow that restarts the vKVM service running on the CIMC. Useful to troubleshoot when launch of vKVM is failing or vKVM becomes unresponsive. Supported with firmware version 4.2(1a) and above.

- **Reset Memory Errors:**

  Resets all memory error counters of all DIMMs to 0. All disabled DIMMs due to prior uncorrectable errors (UECC) will be re-enabled on next boot if they don’t encounter further errors during POST.

- **Clear System Event Logs (SEL):**

  Clears all SEL Logs from the CIMC memory. Useful when they are full. Contain records of hardware events and failures, and other system critical events.

What is VNTag?

- Standardized under 802.1BR reference “Bridge Port Extension”
- Allows creating “virtual cables” between the server adapter and the Fabric Interconnect
- Each “virtual cable” corresponds to an interface from the OS perspective
- Allows complete segmentation of different types of traffic (Management, Ethernet, Storage)
- Tags are applied by the infrastructure and not by the Operating System, to guarantee isolation

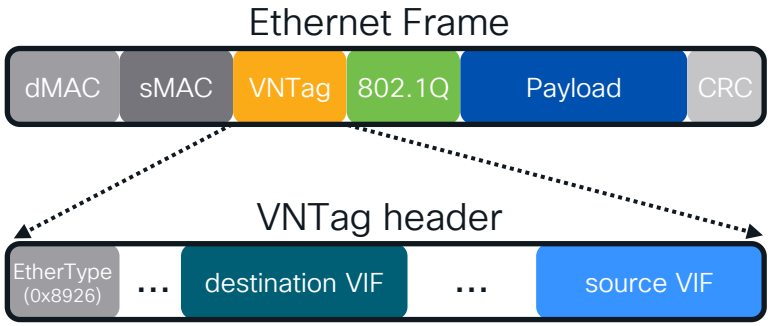

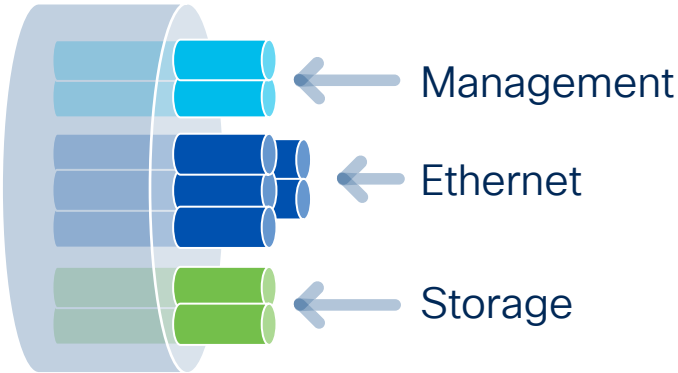

Using VNTag for End-to-End Segmentation in UCS

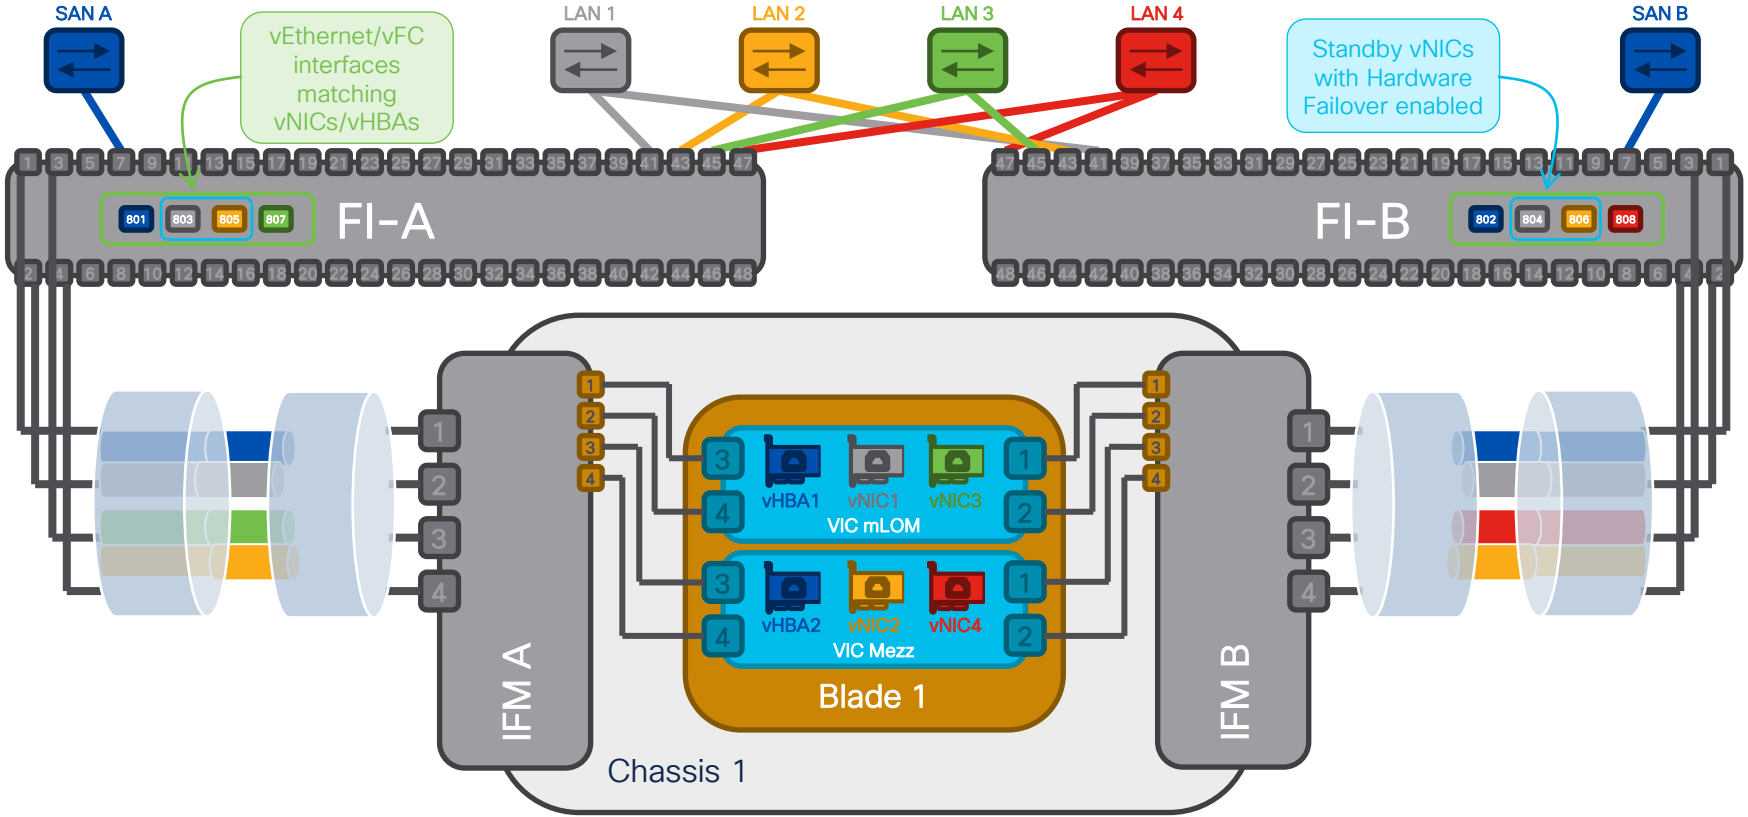

Deploying Server & Chassis Profiles

cisco Live!

Server Profile & Policies

cisco Live!

UCS Server Profile policy structure

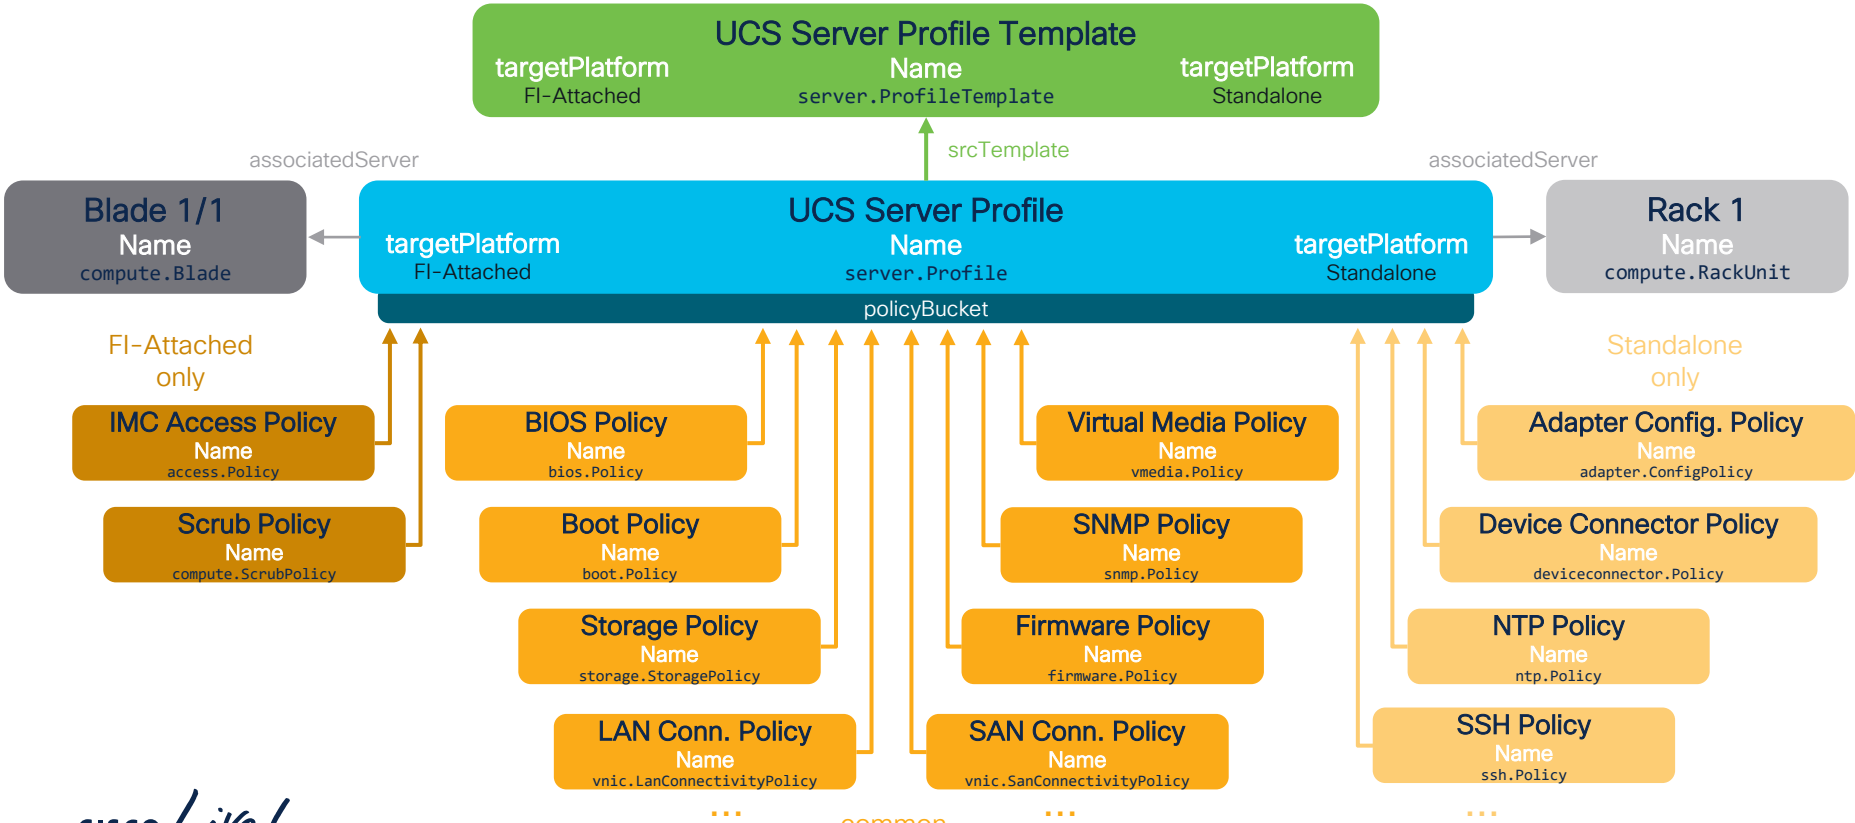

Deploying a Server Profile

- Deploy action always initiated by user
  - Never automatic based on policy change
- Initiates a workflow to deploy policies to Server
  - Validates each policy prior to deployment
- Policy changes requiring server reboot:
  - Changing number/placement of vNICs/vHBAs & QoS settings (LCP/SCP)
  - Changing Firmware, Boot, BIOS, Adapter Config.*, Persistent Memory*
  - Changing some advanced settings of Storage Policy
- Reboot performed as part of profile activation (never automatic)
- Unassign operation cleans up the configuration from the server and resets most policies & settings to factory defaults
  - Except Firmware and storage-related policies. More details here: https://intersight.com/help/saas/configure/servers#clearing_and_resetting_server_configurations

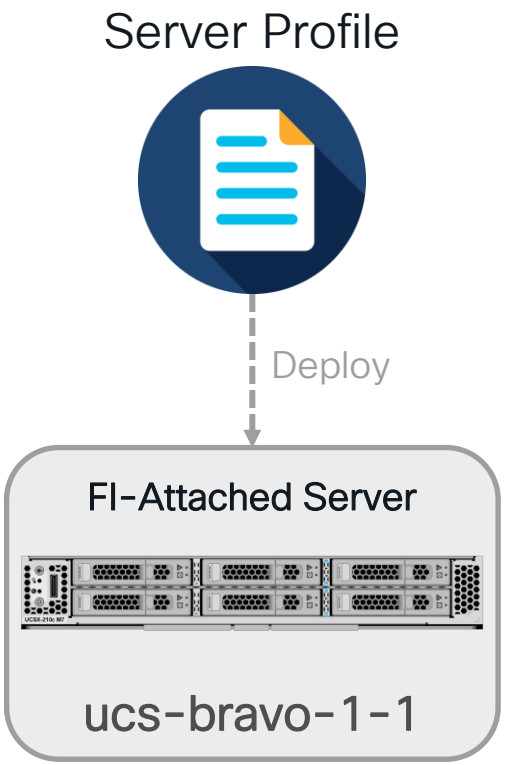

** Only used with ISM, not IMM*

Server Profile Deployment internals

Deploying a vMedia Policy

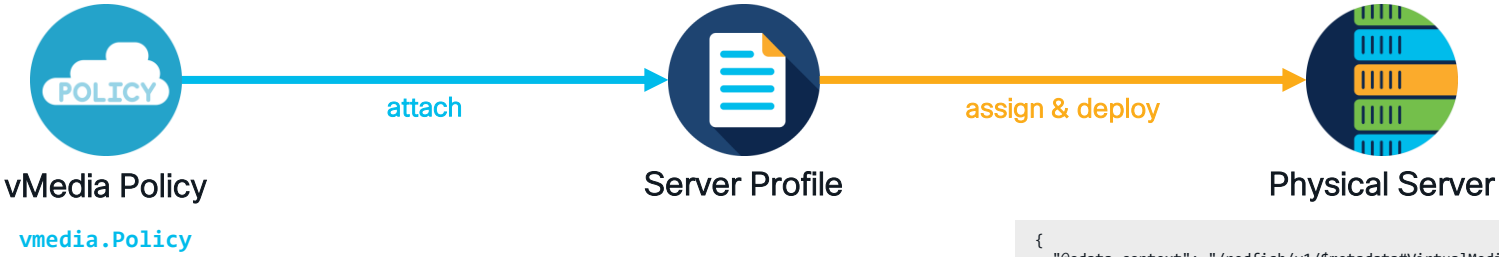

Redfish

```
{
  "@odata.context": "/redfish/v1/$metadata#VirtualMedia.VirtualMedia",
  "@odata.id": "/redfish/v1/Managers/CIMC/VirtualMedia/3",
[snip]
"Id": "3",
  "Image": "https://10.60.10.50/isos/vmware/cisco_images/VMware-ESXi-
7.0.3i-20842708-Custom-Cisco-4.2.2-a.iso?noauto",
  "ImageName": "VMware-ESXi-7.0.3i-20842708-Custom-Cisco-4.2.2-a.iso",
  "Inserted": true,
[snip]
"Name": "CIMC-Mapped vDVD",
  "Oem": {
    "Cisco": {
      "@odata.type": "#CiscoUCSExtensions.v1_0_0.CiscoUCSExtensions",
      "ImageNameVariable": "VMware-ESXi-7.0.3i-20842708-Custom-Cisco-
4.2.2-a.iso",
      "RemapOnEject": null
    }
  },
  "Password": null,
  "Status": {
    "Health": "OK",
    "State": "Enabled"
  },
  "TransferProtocolType": "HTTPS",
[snip]
```

Server Profile Deployment internals

Deploying a LAN/SAN Connectivity Policy (1) – Assigning the Server Profile

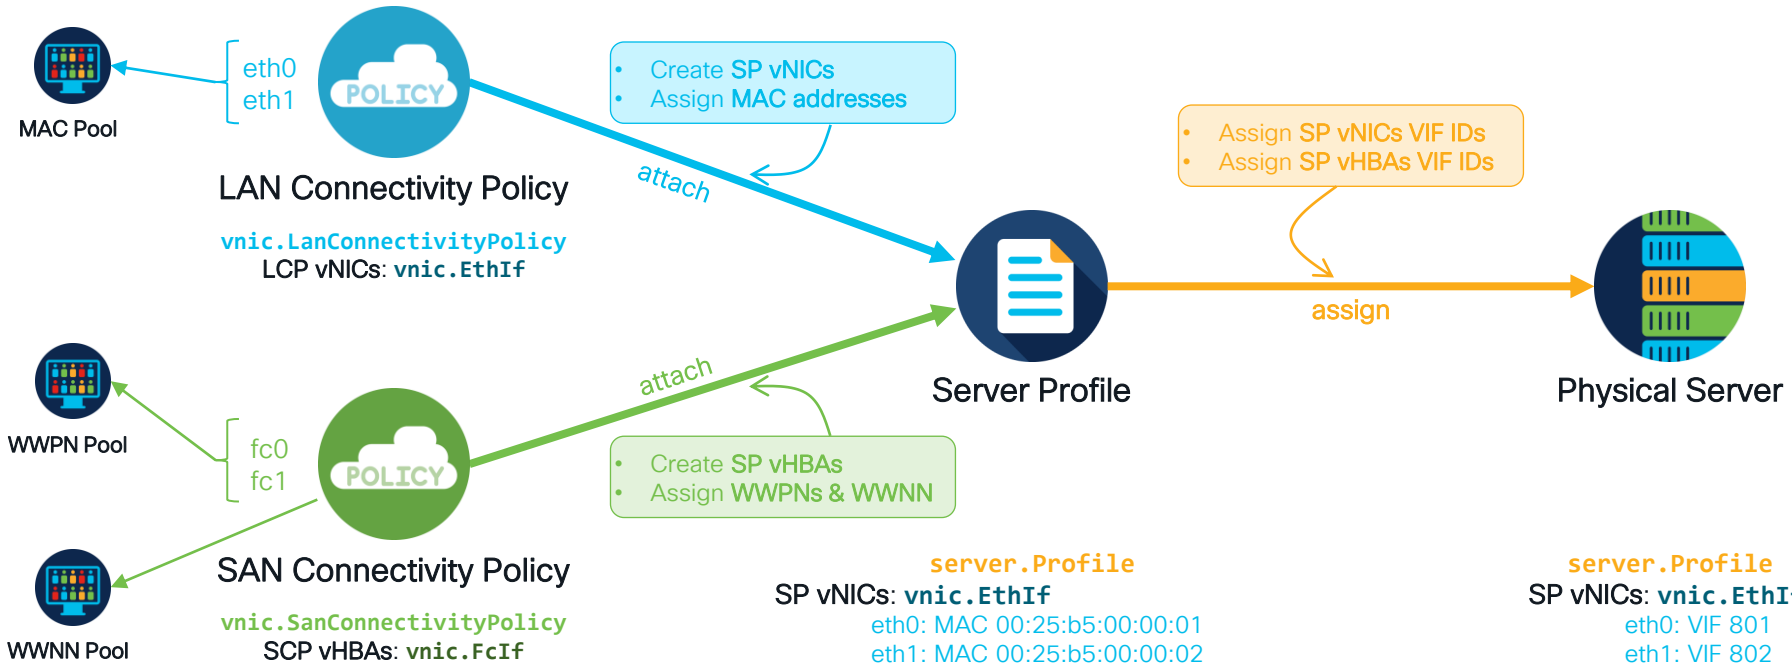

server.Profile
SP vNICs: **vnic.EthIf**
eth0: MAC 00:25:b5:00:00:01
eth1: MAC 00:25:b5:00:00:02
SP vHBAs: **vnic.FcIf**
fc0: WWPN 20:00:00:25:b5:00:00:01
fc1: WWPN 20:00:00:25:b5:00:00:02
WWNN: 20:ff:00:25:b5:00:00:00

server.Profile
SP vNICs: **vnic.EthIf**
eth0: VIF 801
eth1: VIF 802
SP vHBAs: **vnic.FcIf**
fc0: VIF 803
fc1: VIF 804

Server Profile Deployment internals

Deploying a LAN/SAN Connectivity Policy (2) – Deploying the Server Profile

1. Validation checks are performed
   - At server level: compatibility checks, contiguous PCI order, etc.
   - At FI level: pinning checks, VLAN checks, etc.
2. Server-side deployment (**oak**) (using Redfish to CIMC DC)
   - Update vNIC/vHBA config on VIC adapter to match LCP/SCP config
   - Deploy VIC-related policies: Eth/FC Adapter, Eth/FC QoS, iSCSI Boot
   - Asynchronous Intersight inventory update (**dejavu**) for **vnic.EthIf** & **vnic.FcIf** objects
3. FI-side deployment (**gobi**) (using INA messages to FI DC)
   - Create/Update vEthernet/vFC interfaces on NX-OS corresponding to each vNIC/vHBA using VIF ID
   - Deploy NX-OS related policies: Eth Network Group/FC Network, Eth Network Control, Eth/FC QoS, FC Zones
   - Asynchronous Intersight inventory update (**dejavu**) for **network.vEthernet** & **network.vFc** objects

- Unassign unconfigures both VIC and FI configurations

Redfish

```
{
  "@odata.context":
"/redfish/v1/$metadata#NetworkDeviceFunction.NetworkDeviceFunction",
  "@odata.id": "/redfish/v1/Chassis/FCH2712795F/NetworkAdapters/UCSX-V4-
Q25GML_FCH270979C2/NetworkDeviceFunctions/eth0",
[snip]
  "Ethernet": {
    "MACAddress": "00:25:B5:00:00:01",
    "MTUSize": 1500,
[snip]
"Name": ”eth0",
[snip]
  "Oem": {
    "Cisco": {
      "@odata.type": "#CiscoUCSExtensions.v1_0_0.CiscoUCSExtensions",
      "VnicConfiguration": {
[snip]
        "StandByVif": {
          "VifId": 0
[snip]
       "PCIOrder": "0.0",
       "UplinkPort": 1,
       "Vif": {
         "VifCookie": 801,
         "VifId": 801,
[snip]
```

Hardware Failover VIF ID

vNIC VIF ID

NX-OS

```
interface Vethernet801
  description SP demo-1, vNIC eth0, Blade:FCH271278Z5
  no pinning server sticky
  pinning server pinning-failure link-down
  switchport mode trunk
  switchport trunk native vlan 105
  switchport trunk allowed vlan 105
  no hardware vethernet mac filtering per-vlan
  bind interface port-channel1282 channel 801
  service-policy type qos input demo-default
  no shutdown
```

VLANs config

VN-Tag channel

QoS

Inband & Out of Band management (IMC Access)

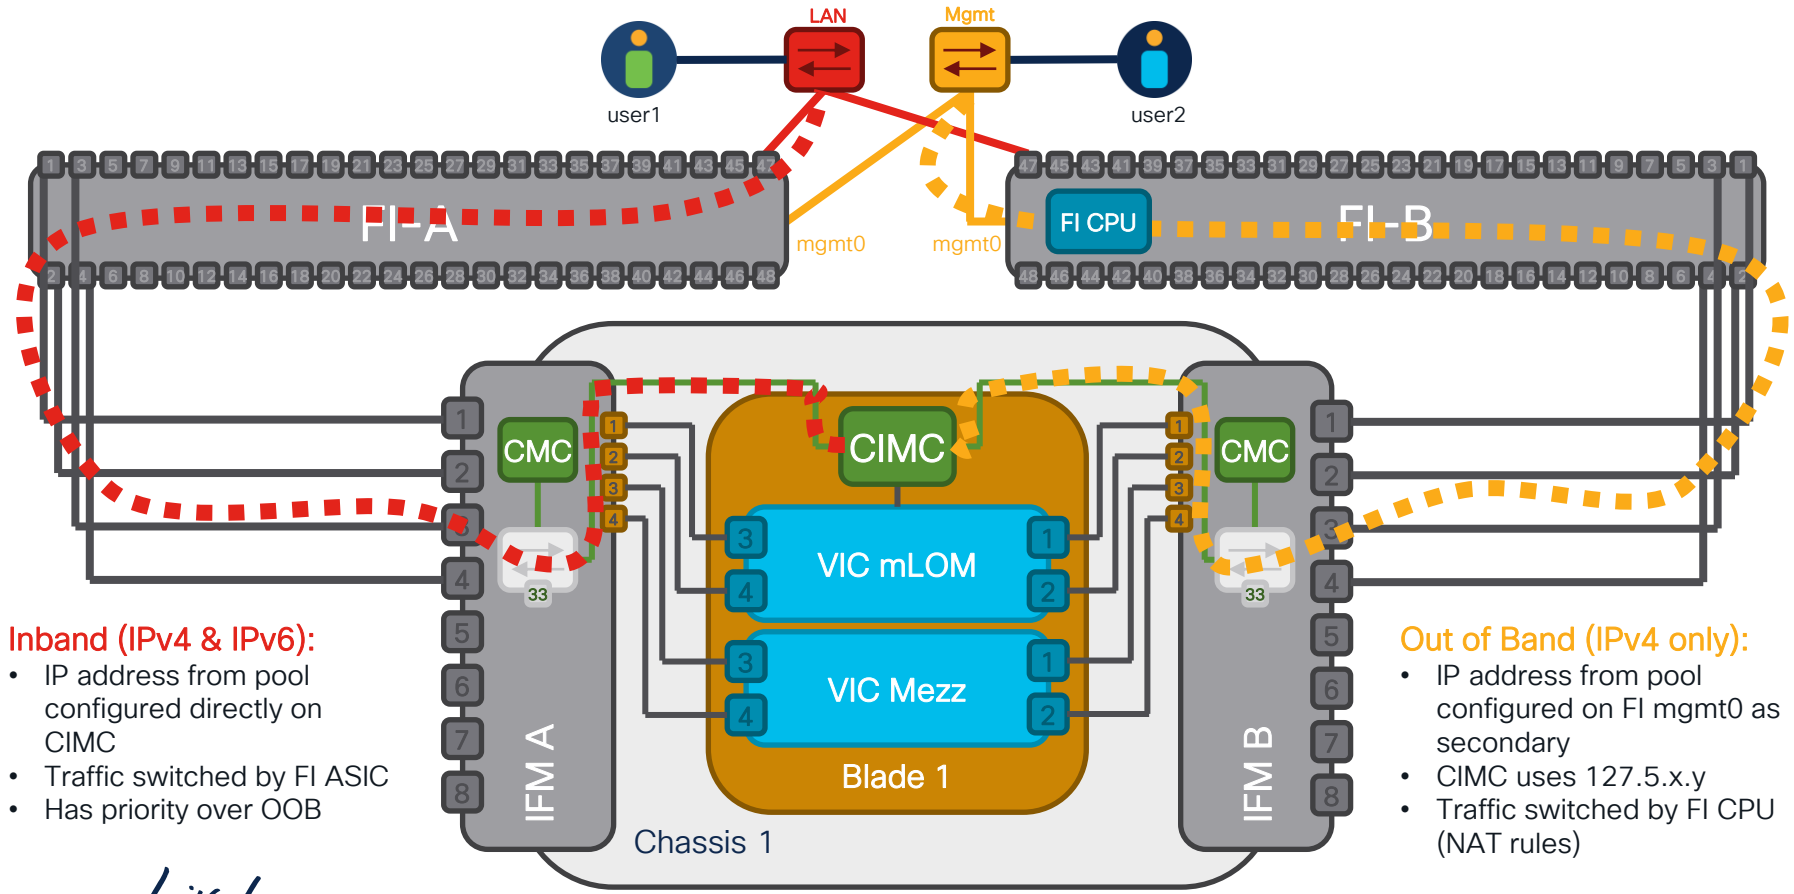

Inband (IPv4 & IPv6):

- IP address from pool configured directly on CIMC
- Traffic switched by FI ASIC
- Has priority over OOB

Out of Band (IPv4 only):

- IP address from pool configured on FI mgmt0 as secondary
- CIMC uses 127.5.x.y
- Traffic switched by FI CPU (NAT rules)

Server Profile change control

1 Assign

- Maps Server Profile to physical server
- “Reserves” hardware resources

2 Deploy

- Pushes configuration to physical server and endpoints (VIC, CIMC, RAID ctrl, etc.)
- Downloads & stages firmwares
- No reboot performed*

3 Activate

- Reboots server to activate deployed changes & firmwares
- Only shows up when disruption is required
- Triggered either by rebooting from Intersight or OS

** Except for Persistent Memory Policy (for ISM)*

Server Profile Pending Changes

- Profile can have different statuses:
  - Not Assigned
  - OK
  - Not Deployed
  - Validating
  - Inconsistent
    (with reason: Pending Changes and/or Out of Sync)
- Pending Changes: difference between the Intersight policy and the last deployed settings
  - Computed when attached policy is changed, or new policy is attached to/detached from profile
  - Profile moves to Validating state while computing
  - Profile moves to Inconsistent if there are pending changes
- Ability to view the pending changes before deploying, with a side-by-side comparison view
- Also available for Domain & Chassis Profiles

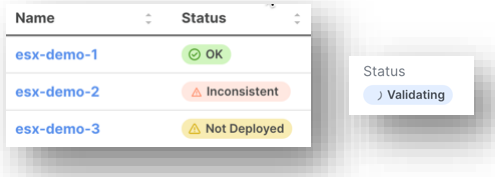

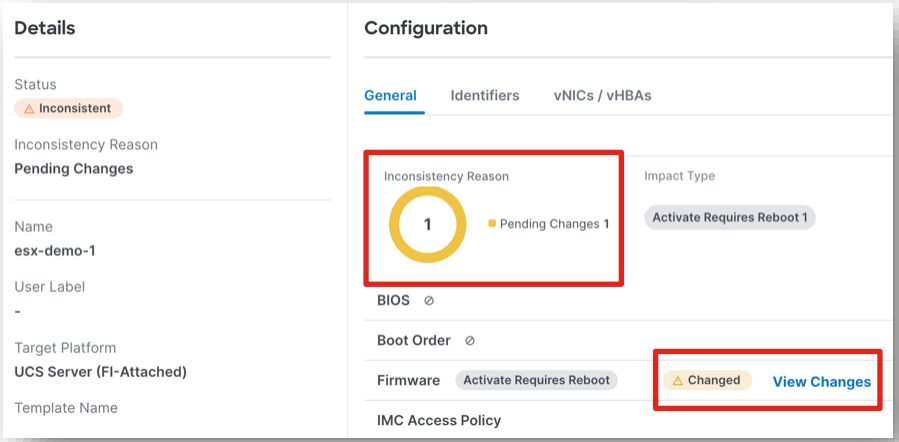

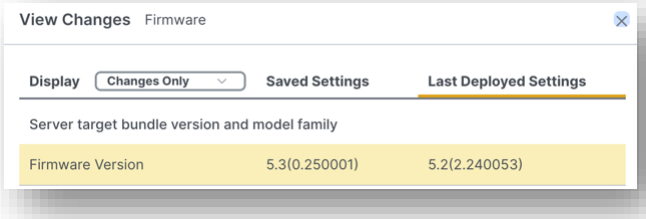

Server Profile Estimate Impact

- Estimates impact of deploying & activating pending changes for each policy
- Impact Types:
  - "Activate Requires Reboot" indicates a reboot requirement to activate the policy
  - "Management Network Outage" indicates a CIMC connectivity loss during deployment
  - "Storage Reconfiguration" indicate a storage configuration change during deployment
- Only available for Server Profiles
  - Chassis Profile deploy is not disruptive
  - Domain Profile displays warning when deploy requires a reboot of FIs

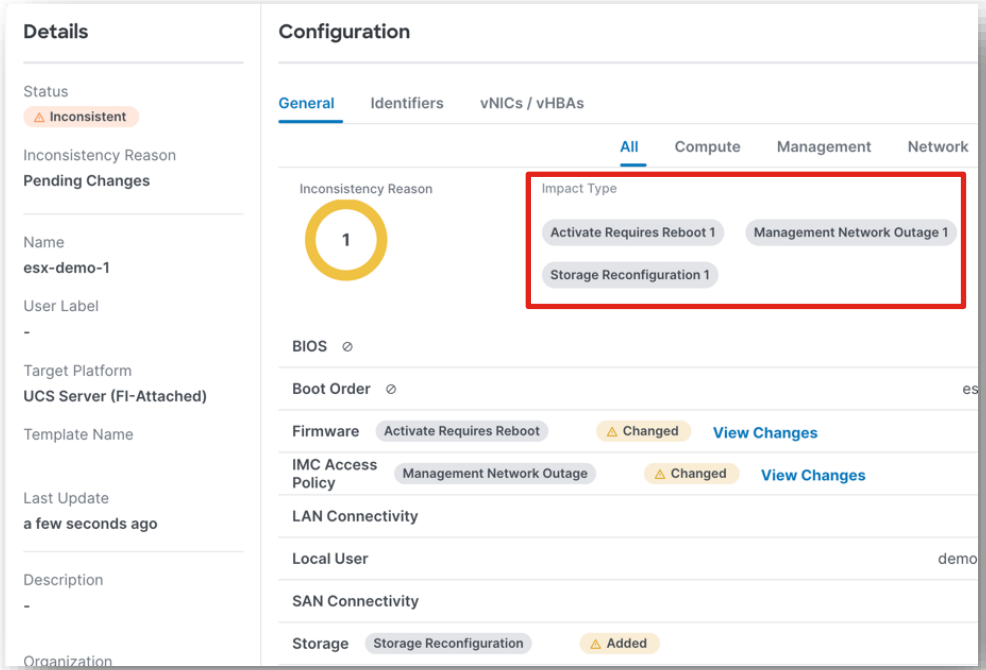

Server Profile Out of Sync

- Reports drift between settings currently configured on the server and last deployed settings from Intersight
- Sequence of operations:
  - After deploy action is complete, Intersight sends CIMC DC a list of policies to track
  - CIMC DC polls config locally using Redfish every 30 minutes and reports back to Intersight if there is a change
  - Intersight compares endpoint settings with last deployed settings (snapshot) for policies that have changed
- Unsupported policies for drift detection: BIOS, Boot, Storage, Certificate Management (future enhancements)
- Only available for Server Profiles

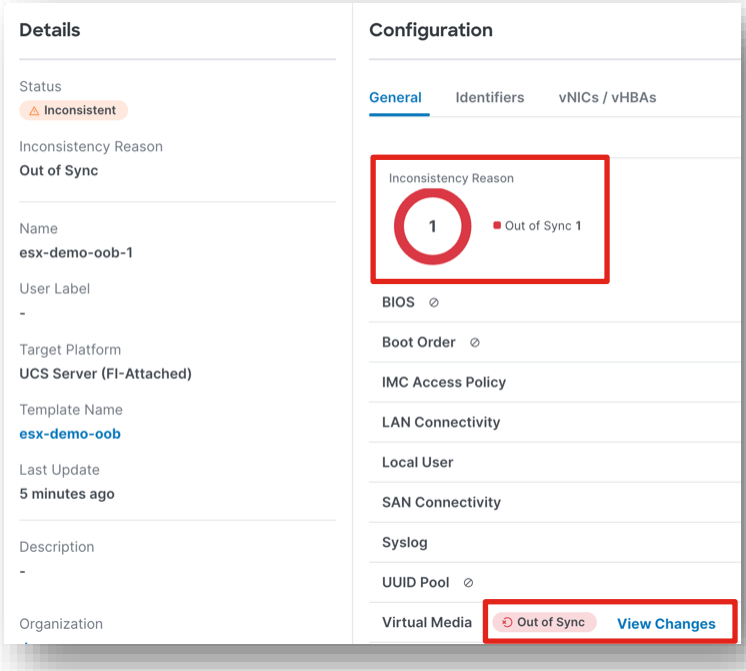

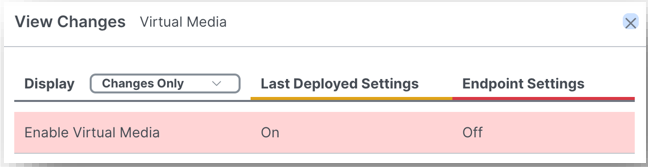

Chassis Profile & Policies

UCS Chassis Profile policy structure

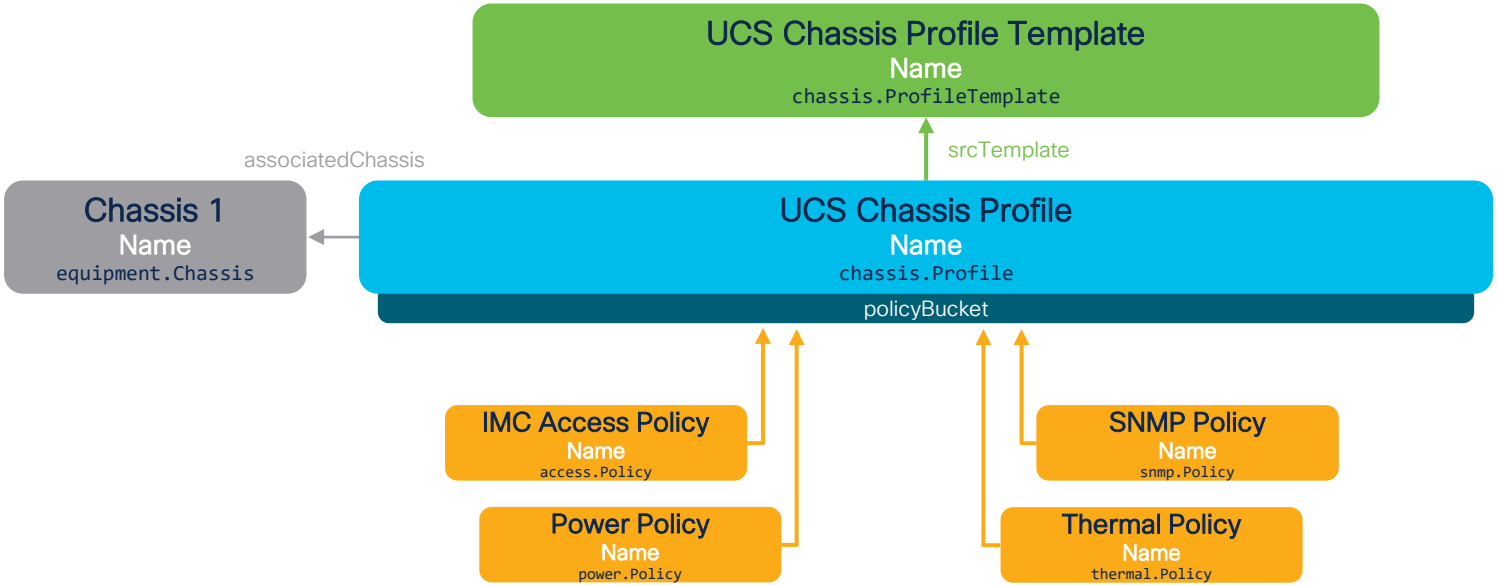

Deploying the Chassis Profile

- Deploy action always initiated by user
  - Never automatic based on policy change
- Initiates a workflow to deploy policies to Chassis
  - Validates each policy prior to deployment
- All policies can be deployed without reboot:
  - Blades can be throttled if not given enough power budget
- IMC Access Policy used to assign IP address to each IOM/IFM in chassis for SNMP
  - Only Inband supported today. Outband support in roadmap
- Unassign operation does not clean up the configuration from the chassis
  - Enhancement in roadmap

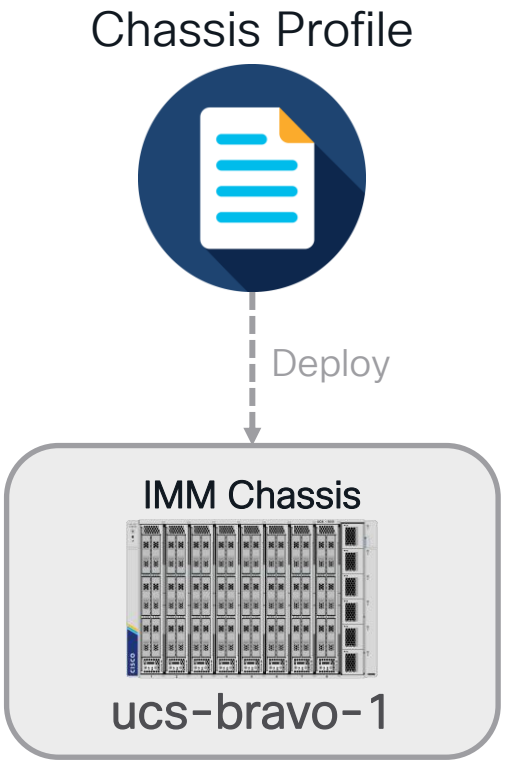

Performing (some) Operations

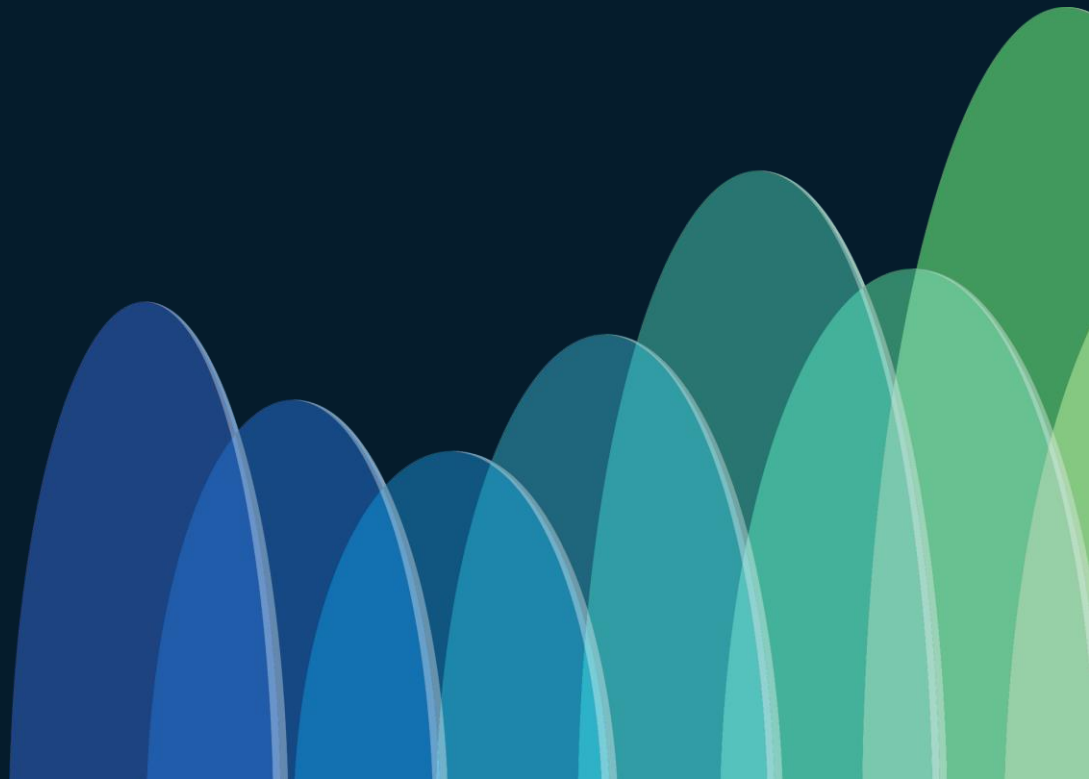

Connect to vKVM

- KVM service running on CIMC next to **jaguar** DC
  - Cisco vKVM with M6+ servers & M5 blades (4.2(2)+)
  - Avocent vKVM on older hardware
- Can be accessed directly via CIMC IP in browser (KVM Direct)
  - Uses Local User Policy to define authentication
  - Requires TCP ports 443 (web UI) & 2068 (vKVM)
- Access via Intersight UI (“Launch vKVM”)
  - Provides SSO capability using OTP (valid for 1min)
  - Requires TCP port 2068 only (web UI content served by Intersight)
- Tunneled vKVM provides access via Intersight
  - vKVM websockets tunneled via Intersight (**propeller**)
  - Requires Advantage license

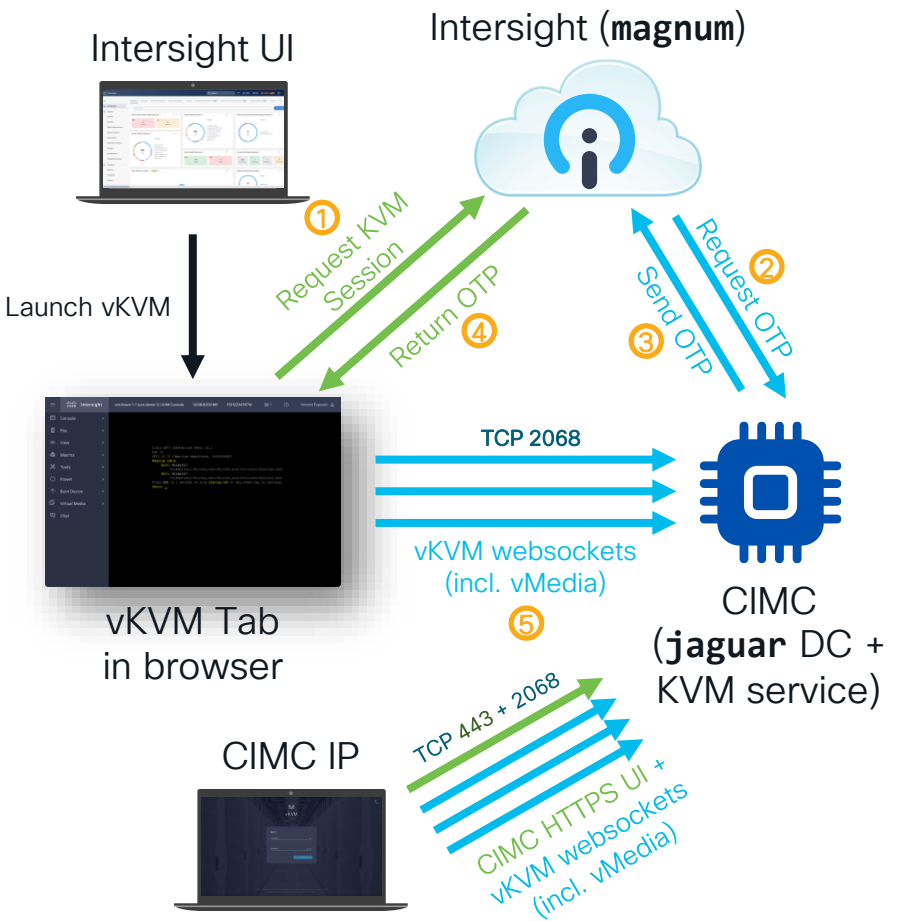

Fabric Evacuation

- Allows evacuating all traffic flowing through a single Fabric Interconnect
  - For IOM/IFM/FEX, shuts down all the Host Interface (HIF) ports
  - For direct-attached rack servers, shuts down all the corresponding server ports (improvement from UCSM)
- If vNICs are using Hardware Failover, traffic is failed over to the standby vNICs on the other FI
- "Force Evacuation" allows enabling Evacuation even if an IOM/IFM/FEX is inoperable, or if a direct-attached rack server port is disabled/down
- Auto-enabled during Infrastructure Firmware upgrades

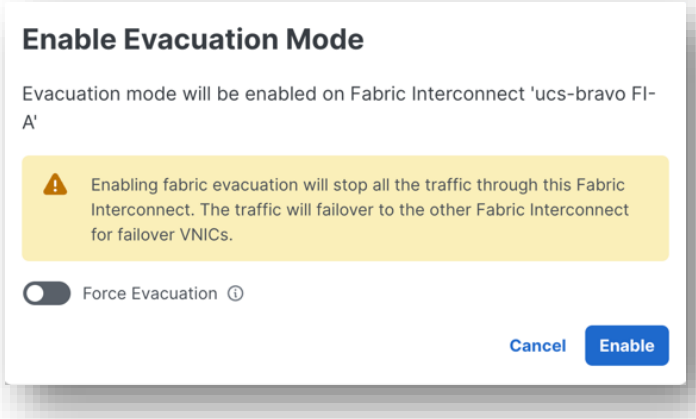


```
ucs-bravo-A(nx-os)# show interface brief
[snip]
--------------------------------------------------------------------------------
Ethernet      VLAN    Type Mode   Status  Reason                   Speed     Port
Interface                                                                    Ch #
--------------------------------------------------------------------------------
Eth1/17       1       eth  fabric up      none                     25G(D) 1025
Eth1/18       1       eth  fabric up      none                     25G(D) 1025
[snip]
Eth1/33       1       eth  vntag  down    Administratively down    auto(D) 1280
[snip]
Eth1/49       1       eth  trunk  up      none                     100G(D) --
Eth1/50       1       eth  trunk  up      none                     100G(D) --
[snip]
Eth1/1/1      1       eth  vntag  down    Administratively down    auto(D) 1286
Eth1/1/2      1       eth  vntag  down    Administratively down    auto(D) 1286
Eth1/1/3      1       eth  vntag  down    Administratively down    auto(D) 1287
Eth1/1/4      1       eth  vntag  down    Administratively down    auto(D) 1287
[snip]
```

IFM NIF

rack server

uplinks

IFM HIF

Upgrading infrastructure firmwares (1)

- Action performed on Fabric Interconnect
  - Upgrades entire domain by default
  - Single Fabric Interconnect Upgrade option
  - Evacuate FI traffic option (enabled by default)
- Tracks IO Path Connectivity during upgrade
  - Ensures all data path alarms are cleared prior to allowing user-ack
  - 60min timer for alarms to clear. Once timer expires, user-ack will be allowed regardless of the data path alarms status
  - Can be disabled in Advanced Options
- Platform-specific infra firmware bundle
  - System image: FI NX-OS, `jaguar` DC & `dove` management package
  - IOM/IFM/FEX images: NX-OS, CMC and `jaguar` DC (per model image)
  - PSU images: Chassis PSU firmwares

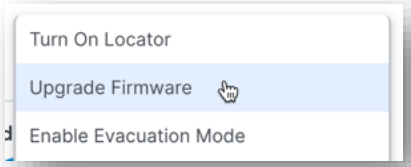

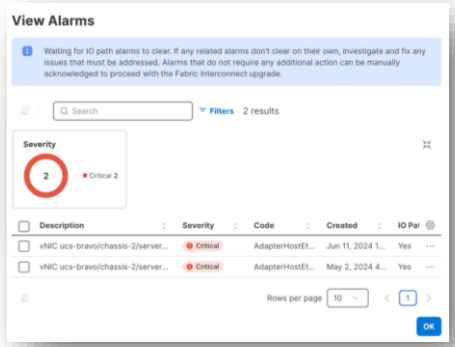

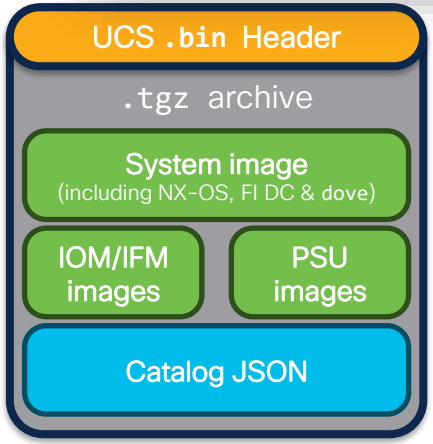

Infrastructure firmware bundle

Upgrading infrastructure firmwares (2)

Sequence of Operations

1. Download infra bundle to FI-A cache if not present (and sync to FI-B)
2. Upgrade B-side infrastructure:
   - Upgrade all FI-B IOM/IFM/FEX (without activation)
   - Upgrade FI-B NX-OS (without activation)
   - Evacuate FI-B
   - User-ack to reboot FI-B
   - Activate FI-B and IOM/IFM/FEX images & Reboot FI-B
3. Wait for B-side data path alarms to clear & user-ack from Intersight
4. Upgrade A-side infrastructure:
   - Upgrade all FI-A IOM/IFM/FEX (without activation)
   - Upgrade FI-A NX-OS (without activation)
   - Evacuate FI-A
   - Activate FI-A and IOM/IFM/FEX images & Reboot FI-A

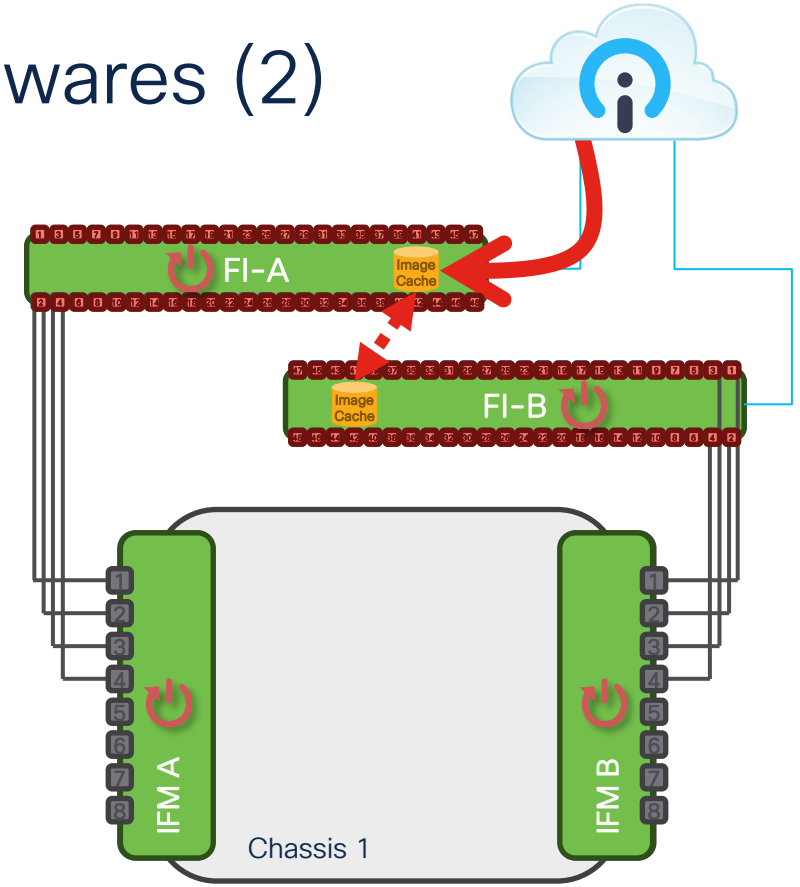

Upgrading server firmwares (1)

- Can be triggered via user action or Firmware Policy deploy/activate
- Platform-specific server firmware bundle (HSU – Host Server Utility) containing:
  - OOB components: CIMC, BIOS, VIC
  - Non-OOB components: RAID, GPU, drives, etc.
- Sequence of operations:
  1. Download HSU bundle to FI-A cache if not present (and sync to FI-B)
  2. Upgrade OOB components:
     - CIMC can download OOB components from FI cache directly
     - Upgrade all VIC adapters, CIMC and BIOS to backup bank and activate on next reboot
     - Can be monitored on CIMC via `sw_update_status <bmc/bios>`
- If action did not include "Reboot Immediately", or Profile Activation is not triggered, firmware upgrade is paused until next reboot
  - Reboot can be performed from Intersight or from OS

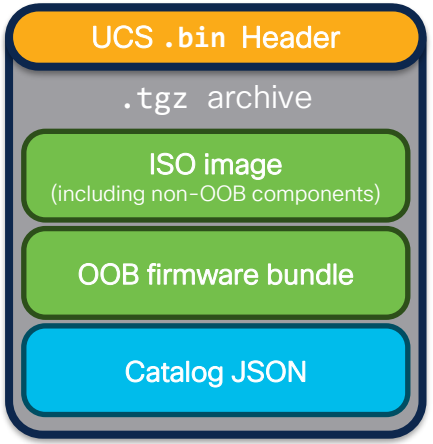

HSU server firmware bundle

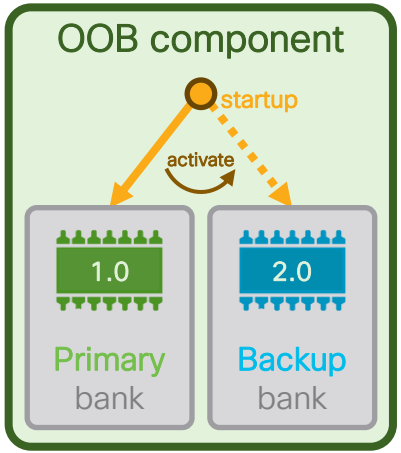

UCS OOB component upgrade

Upgrading server firmwares (2)

- If “Reboot Immediately to Activate” is selected (action), or Profile is activated (policy), or server is rebooted:
  3. HSU agent (running on CIMC) configures HTTPS vMedia Mount with one-time boot of ISO image from HSU bundle in FI cache
  4. After ISO boot, HSU agent performs inventory scan & updates all non-OOB components
  5. HSU agent reports back progress to Intersight via Redfish (status polled every minute)
  6. Reboot host to OS
- Triggers Blade/Rack Discovery (“deep”) workflow to update Intersight inventory
- HSU progress can be monitored via vKVM
- UCS Firmware Equivalency Matrix: https://www.cisco.com/c/dam/en/us/td/docs/unified_computing/ucs/c/sw/UCS-Equivalency-Matrix/index.html

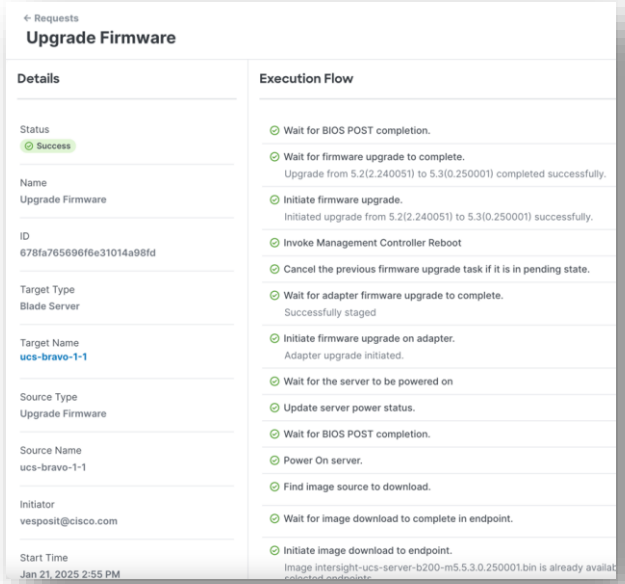

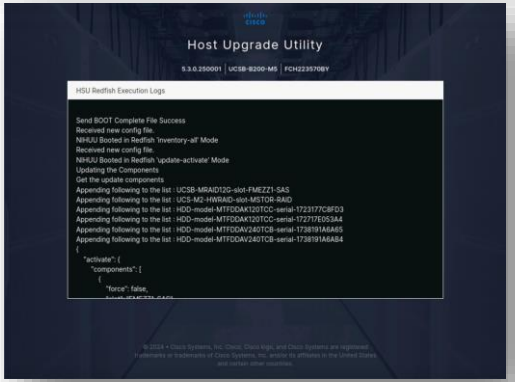

Managing FI firmware cache

- All firmwares (infrastructure & server bundles) used to perform a firmware upgrade are automatically stored on FI cache
  - Downloaded from intersight.com for SaaS & CVA
  - Downloaded from Software Repository for PVA
  - Synchronized between both FIs
- If cache becomes full, it will prevent firmware to be downloaded
- Monitor cache space utilization under the FI inventory
- To clear cache space, use the `clear-firmware-cache` CLI command, specifying a firmware bundle ID to remove (use `list-firmware-cache` to list bundle IDs)
  - Perform clear on **both** Fabric Interconnects
- Automatic cache space management in roadmap

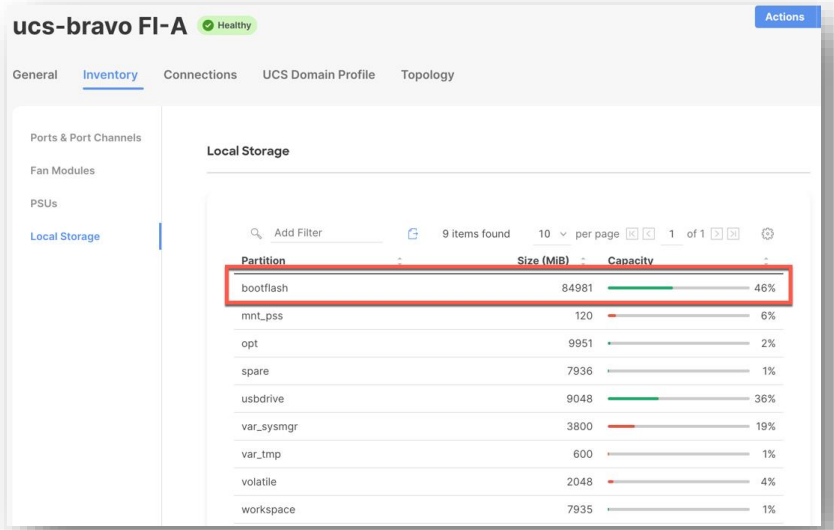


```
ucs-bravo-A# list-firmware-cache
DOWNLOADED DATE     IDENTIFIER                        IMAGE NAME
2023-09-13          650216eb65676131019d57ad          ucs-intersight-infra-4gfi.4.3.2.230117.bin
2023-10-30          653fbea3656761310100d082          intersight-ucs-server-b200-m5.4.2.3f.bin
2023-11-02          6543c63e656761310112aeba          intersight-ucs-server-b200-m5.5.1.0.230054.bin
2023-11-19          655a7771656761310121512d          intersight-ucs-server-210c-m6.5.2.0.230092.bin
2023-11-19          655a77f56567613101215368          intersight-ucs-server-b200-m5.5.2.0.230100.bin
2023-11-19          655a8bd9656761310121a9aa          intersight-ucs-server-c220m5.4.3.2.230270.bin
2023-11-20          655bd6bd6567613101276637          ucs-intersight-infra-4gfi.4.3.2.230129.bin
2023-11-22          655e0f9265676130057f6761          intersight-ucs-server-b200-m5.5.1.0.230073.bin
2023-12-11          6576dd556567613101d09d47          intersight-ucs-server-b200-m5.4.2.3b.bin
2024-01-24          65b1355d65676130059f01b7          intersight-ucs-server-b200-m5.5.1.0.230073.bin
```

```
ucs-bravo-A# clear-firmware-cache 6543c63e656761310112aeba
6543c63e656761310112aeba
The Intersight cache will be cleared.
Are you sure? Enter 'y' to continue:y
Cache cleared
```

For your reference

Hardware Compatibility List compliance checks

- Leverages UCS-Tools (VIB for ESXi)
  - Executed at each ESXi boot
  - Gathers inventory & driver details
  - Sends info to CIMC over IPMI (local)
- Intersight automatically validates compliance
  - Not Listed: firmwares/drivers installed not in HCL database
  - Validated: all firmwares/drivers compliant
- Open-Source equivalents for Win & Linux:
  - https://github.com/CiscoDevNet/intersight-powershell-utils/tree/main/os-discovery-agent-windows
  - https://github.com/CiscoDevNet/intersight-python-utils/tree/main/os-discovery-agent-linux

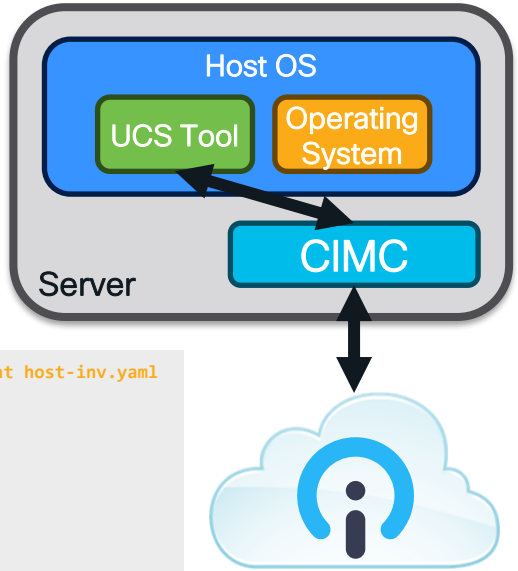


```
[root@esx-prod-1:/opt/ucs_tool_esxi] cat host-inv.yaml
annotations:
[snip]
 -kv:
  key: os.vendor
  value: VMware ESXi
 -kv:
  key: os.kernelVersionString
  value: 8.0.3
 [snip]
 -kv:
  key: os.driver.0.name
  value: nenic
 -kv:
  key: os.driver.0.description
  value: Cisco Systems Inc Cisco VIC Ethernet NIC
 -kv:
  key: os.driver.0.version
  value: 2.0.11.0-1OEM.800.1.0.20143090
 [snip]
```

Key Takeaways

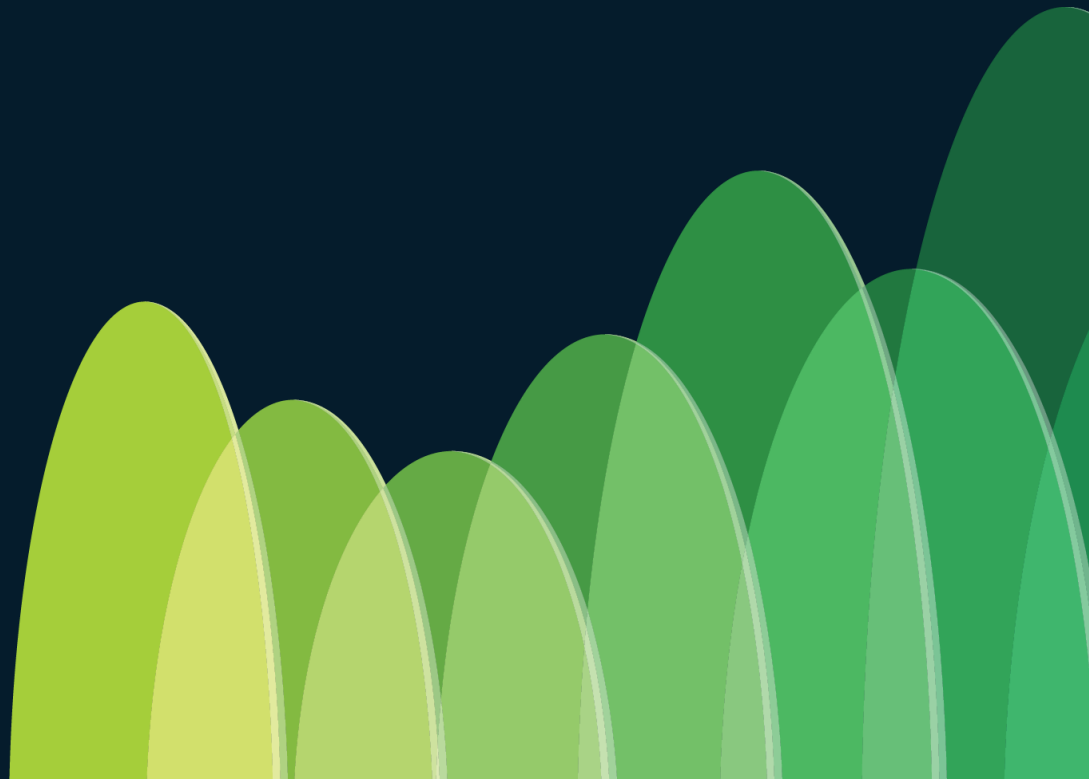

Key IMM characteristics

- Modernized UCS platform leveraging micro-services and open standards (OpenAPI, Redfish)
- Fully automated hardware discovery & inventory for blades & racks
- Flexible advanced connectivity with VIC and VNTag with secure traffic segregation
- Policies & Profiles everywhere (Servers, Chassis, Domain)
- Granular change control (Assign, Deploy, Activate)
- Simple operations

Webex App

Questions?

Use the Webex app to chat with the speaker after the session

How

1. Find this session in the Cisco Events mobile app
2. Click “Join the Discussion”
3. Install the Webex app or go directly to the Webex space
4. Enter messages/questions in the Webex space

Webex spaces will be moderated by the speaker until February 28, 2025.

Fill Out Your Session Surveys

Participants who fill out a minimum of 4 session surveys and the overall event survey will get a unique Cisco Live t-shirt.

(from 11:30 on Thursday, while supplies last)

All surveys can be taken in the Cisco Events mobile app or by logging in to the Session Catalog and clicking the ‘Participant Dashboard’

Content Catalog

Continue your education

- IMM Transition Interactive Breakout: IBOCOM-2301
- Cisco Compute YouTube Channel youtube.com/@CiscoCompute
- Visit us at the Compute Booth
- Attend our Capture the Flag, and Walk-in Labs
- Visit the On-Demand Library for more sessions at ciscolive.com/on-demand. Sessions from this event will be available from March 3.

Cisco Compute YouTube Channel

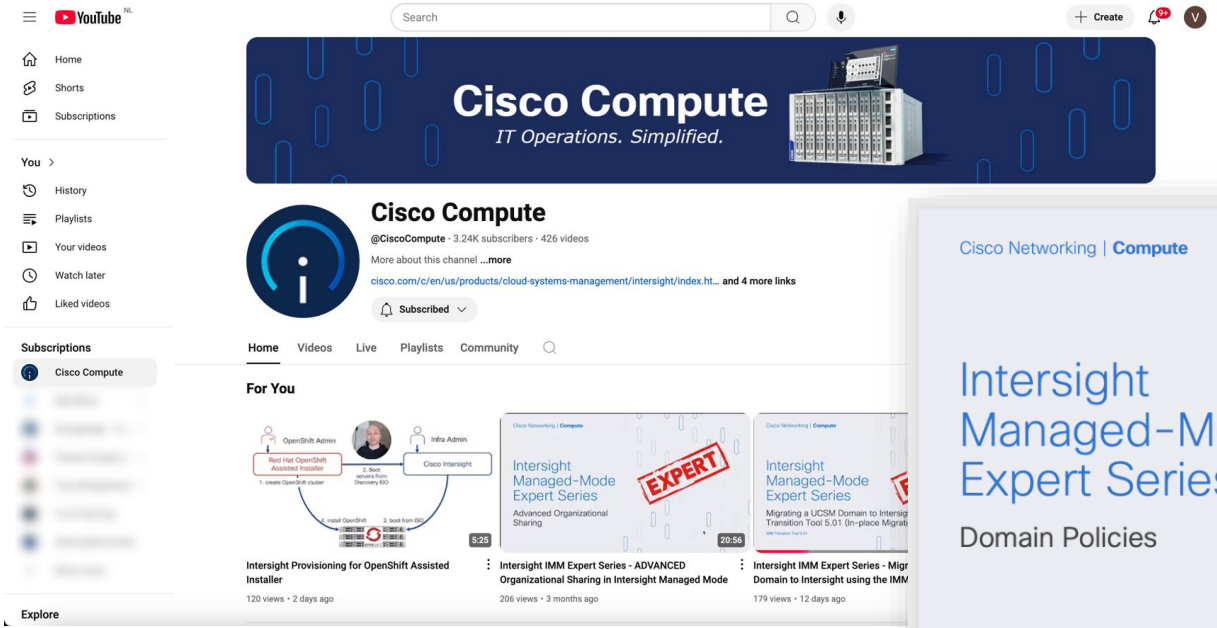

IMM Expert Series

Detailed IMM Training Videos

https://www.youtube.com/@CiscoCompute


cisco
Thank you
cisco Live!

cisco Live!

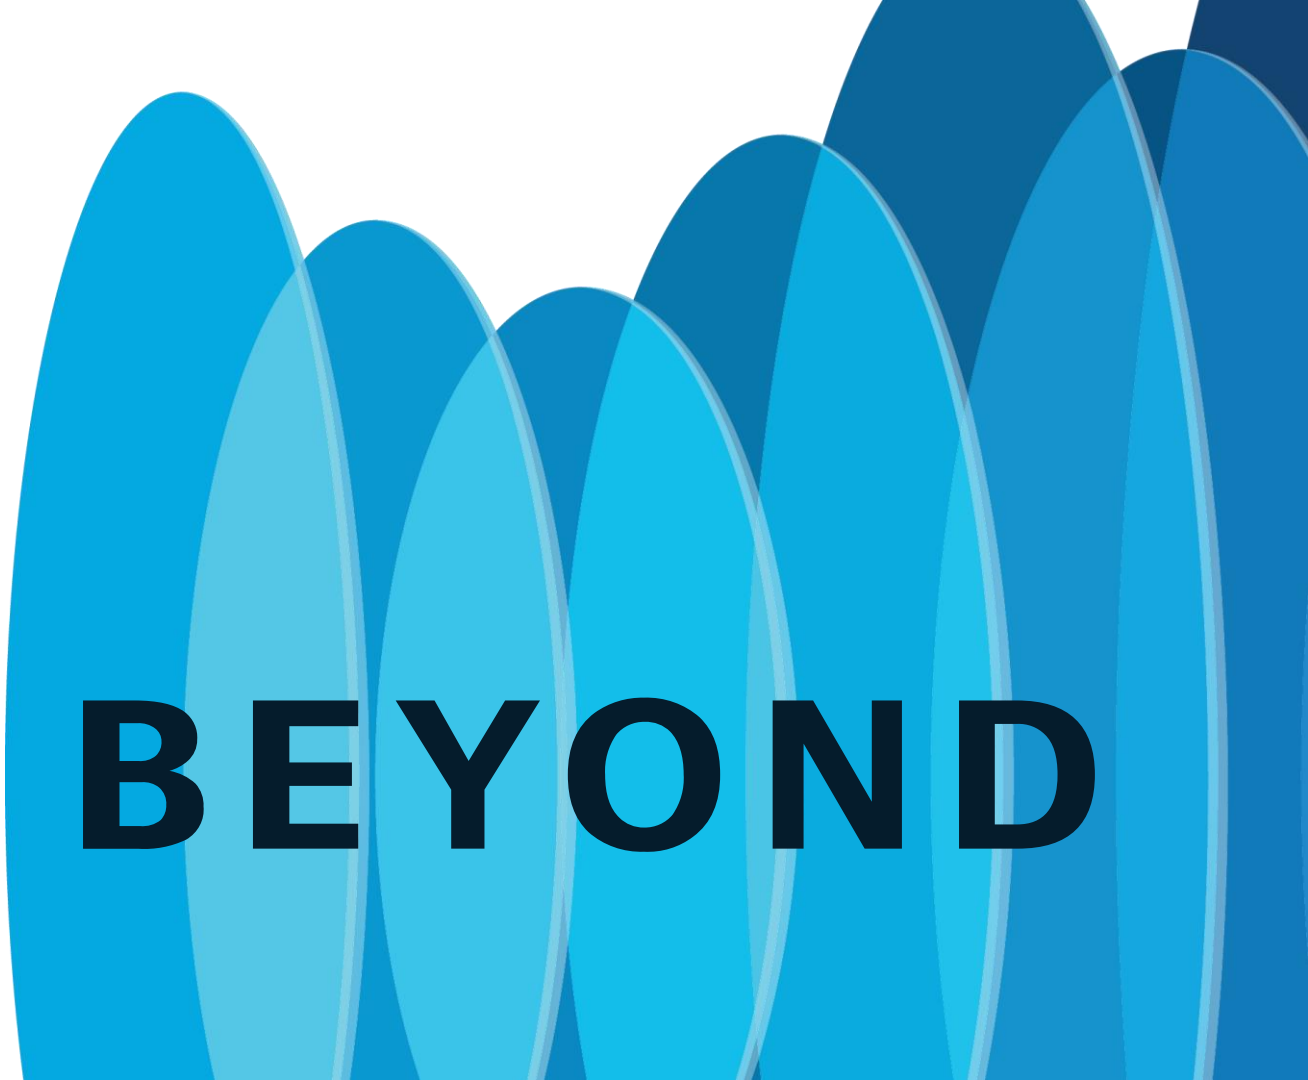

GO BEYOND