# Cisco Catalyst Instant Access Solution

## What You Will Learn

Cisco Catalyst[®] Instant Access creates a single network touch point and a single point of configuration across distribution and access layer switches, dramatically simplifying design, deployment, and operations for enterprise campus networks. This paper discusses the Cisco[®] Catalyst Instant Access Solution's architecture, components, packet walks, and value proposition.

## Overview

Cisco Catalyst Instant Access enables the merging of physical distribution and access layer switches into a single logical entity with a single point of configuration, management, and troubleshooting. The solution simplifies enterprise campus networks by bringing in provisioning and operational simplicity.

Benefits of Cisco Catalyst Instant Access include:

- Single point of configuration and management
- Single software image across distribution and access layers
- "Plug and play" provisioning of access switches
- Agile infrastructure at the access layer, with feature and hardware consistency
- Automatic uplink configuration at the access layer
- Automatic image provisioning of access switches
- Rich and consistent Catalyst 6500/6800 Series feature set across distribution and access layers

Figure 1 depicts a single touch point for a 21 access switch (1000-port) distribution block.

**Figure 1.**    Single Logical Switch



With Cisco IOS[®] Software Release 15.1(2)SY, the Instant Access solution supports:

- 1008 host ports across 21 Instant Access clients
- Stacking of up to three clients

With Cisco IOS Software Release 15.2(1)SY1, the Instant Access solution supports:

- 1536 host ports on Supervisor 2T-based systems on a modular chassis (1200 ports with release 15.2(1)SY)
- Up to 2016 host ports on the Cisco Catalyst 6880-X Switch
- Stacking of up to five clients

Consider the topology outlined in Figure 2. It includes a 4032-port campus network with four distribution blocks, each consisting of 1008 ports (21 access switches of 48 ports each) with a Cisco Virtual Switching System (VSS) pair at the distribution, and stacking technology at the access layer.

This campus requires:

- 29 devices for configuration management
- 29 devices for image management
- 48 trunks and port-channel configurations on access switches
- 29 separate configurations, including Simple Network Management Protocol (SNMP), Network Time Protocol (NTP), TACACS/RADIUS, VLAN database, management IP, gateway, and host name

As shown in Figure 3, with Cisco Catalyst Instant Access, the same 4032-port campus would require only:

- Five total devices to manage
- No image management at access switches
- No uplink trunk configuration on access switches
- Five separate configurations for SNMP, NTP, TACACS, VLAN DB, management IP, and hostname

**Figure 2.**   Traditional Deployment                            **Figure 3.**   Instant Access Deployment



### System Components

The Instant Access solution has two components: the Instant Access parent and the Instant Access client (Figure 4).

**Instant Access parent:** The Instant Access parent switch comprises a pair of Cisco Catalyst 6500-E or 6807-XL Series chassis with Supervisor 2T configured in VSS or VSS Quad-Sup SSO[1] mode and a WS-6904 40G/10G line card configured in 10-Gb mode. Details on configuring in VSS and VSS Quad-Sup SSO mode can be found here. Details on 40-Gb line cards operating in 10-Gb mode can be found here.

[1] Certain deployments may not have a VSS pair at the distribution layer. In such cases, a single Cisco Catalyst 6500 or 6800 Series switch can be used. The switch needs to be configured in VSS mode because the Instant Access solution treats each Instant Access client as a remote line card and uses a VSS infrastructure to enable this remote-line-card-like capability. It is not recommended to deploy Instant Access with a single switch at the distribution layer; however, if configured, it is recommended to have two supervisors in the chassis in case of failure of one supervisor engine. However, note that since Instant Access requires VSS mode, if a single chassis is used with dual supervisors and the active supervisor goes down, the entire system reloads (no RPR or SSO is maintained between the supervisors). (A single chassis with dual supervisors without Instant Access does not require VSS mode, and therefore can maintain Route Processor Redundancy (RPR) or stateful switchover (SSO) between the supervisors.).

Instant Access parent functionality is also supported on a Cisco Catalyst 6880-X or 6880-X-LE Switch VSS pair. Instant Access is supported on all ports on the baseboard and on the port cards of this chassis. With the Release of 15.2(2)SY, the newer Cisco Catalyst 6840-X series switches also support Instant Access parent functionality on the onboard 10 Gigabit Ethernet (10-GE) and 1-GE ports, across all the four models of switches.

With Release 15.2(1)SY, the latest 10-GE line cards in the modular portfolio include Instant Access parent functionality on Cisco Catalyst 6500-E Series and 6807-XL Chassis. A complete list of hardware supporting the Cisco Catalyst Instant Access parent functionality follows.

- Modular chassis: Cisco Catalyst 6500-E Series chassis or 6807-XL chassis

  The modular chassis requires Supervisor Engine 2T (VS-S2T-10G or VS-S2T-10G-XL) along with any of the following line cards:

  ◦ WS-X6904-40G-2T, WS-X6904-40G-2TXL (Instant Access ports in 10-Gb mode)

  ◦ C6800-32P10G, C6800-32P10G-XL

  ◦ C6800-16P10G, C6800-16P10G-XL

  ◦ C6800-8P10G, C6800-8P10G-XL

- Fixed Chassis: All ports of Catalyst C6880-X and C6880-X-LE switches, and all models of the Catalyst 6840-X Switch

**Instant Access client:** The Instant Access client is a Cisco Catalyst 6800ia Series switch operating exclusively in client mode with a Cisco Catalyst 6500-E or 6800 Series switch at the distribution layer. The 6800ia switch is not intended to be used at the aggregation layer, and it is best practice to connect the 6800ia to end hosts or compact switches downstream. The Cisco Catalyst IOS Release 15.2(1)SY adds support for a compact switch in the Catalyst 3560-CX family as an Instant Access client, which can operate either as a standalone switch or as an Instant Access client. 15.2(1)SY1 adds support for a second compact switch as a client, which also belongs to the Catalyst 3560CX family and supports mGig (or Multirate Gigabit Ethernet) connectivity.

The 6800ia Instant Access client supports 48 10/100/1000 interfaces and two 10 Gbps uplink or fabric interface ports. The high-level features and capabilities of the Instant Access client are:

- 48 10/100/1000 BASE-T host ports with Power over Ethernet+ (PoE+) or non-PoE options
- Two 10-Gbps uplink ports
- 740 Watts PoE power:
  ◦ Full PoE (15 Watts) across all 48 ports
  ◦ Full PoE+ (30 Watts) across any 24 ports
- Stackable up to five clients
- 80 Gbps of bidirectional stack bandwidth

- Operates in Instant Access client mode only with centralized packet switching on the Instant Access parent
- A separate SKU with a redundant power supply is available

The Cisco Catalyst 6800ia families of switches have differing power configurations. The first two models, C6800IA-48TD (data-only), and C6800IA-48FPD (PoE/PoE+), support a single, built-in power supply and fan. The power redundancy for these two models is supported by an external Cisco Redundant Power System (RPS). A third model, C6800IA-48FPDR (PoE/PoE+), supports two redundant removable power supplies, and each power supply has a power budget of 1025 Watts of total system power. The 1025-Watt power supply has an inline power budget of 740 Watts.

More details are available here.

**Figure 4.** Cisco Catalyst Instant Access Components
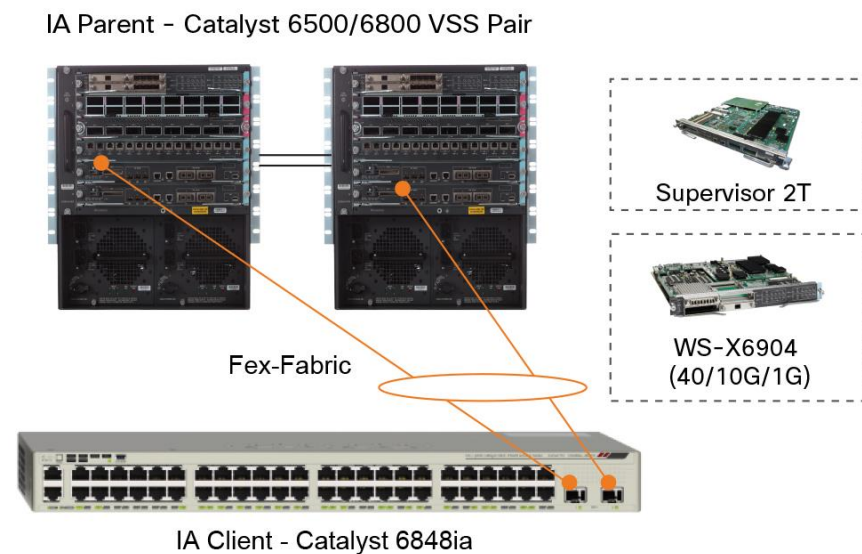


In addition to the parent and client, a **fex-fabric** link between the Instant Access parent and client supports short-reach, long-reach multimode, long-reach, and extended-reach optics with Cisco 10GBase SFP+ across fabric links. For more details, click here.

As mentioned above, with Cisco Catalyst IOS Release 15.2(1)SY, a new Instant Access client has been added to the portfolio. It is the Cisco Catalyst 3560-CX Series compact switch (Figure 5; SKU: WS-C3560CX-12PD-S), which is able to function as an Instant Access client connected to the Catalyst 6500/6800 parent switch.

**Figure 5.** 3560-CX Series: New Instant Access Compact Client

The high-level features and capabilities of this Instant Access client include:

- Twelve 10/100/1000 BASE-T host ports with PoE
- Two 10-Gbps SFP+ or 1-Gbps SFP uplink ports (used for Instant Access)
- Two 10/100/1000 BASE-T uplink ports (not used for Instant Access)
- 240 Watts of PoE power:
  - Full PoE (15 Watts) across all 12 ports
  - Full PoE+ (30 Watts) across any 8 ports
- Does not support stacking
- Fanless switch with a single built-in power supply
- Operates in both Instant Access mode and standalone mode

The 12-port 3560-CX switch is available in three SKUs: WS-C3560CX-12TC-S, WS-C3560CX-12PC-S, and WS-C3560CX-12PD-S. Of these products, only the WS-C3560CX-12PD-S is supported as an Instant Access client.

More details are available here.

With a subsequent release, a second compact switch, WS-C3560CX-8XPD-S, is supported as an Instant Access client. More details can be found here.

The Instant Access solution supports a mix of both 6800ia and 3560-CX switches as clients from the same parent switch. Components are shown in Figure 6.

**Note:** 3560-CX compact switches are not supported in Catalyst 6840 family as Instant Access Client

**Figure 6.** Instant Access Components
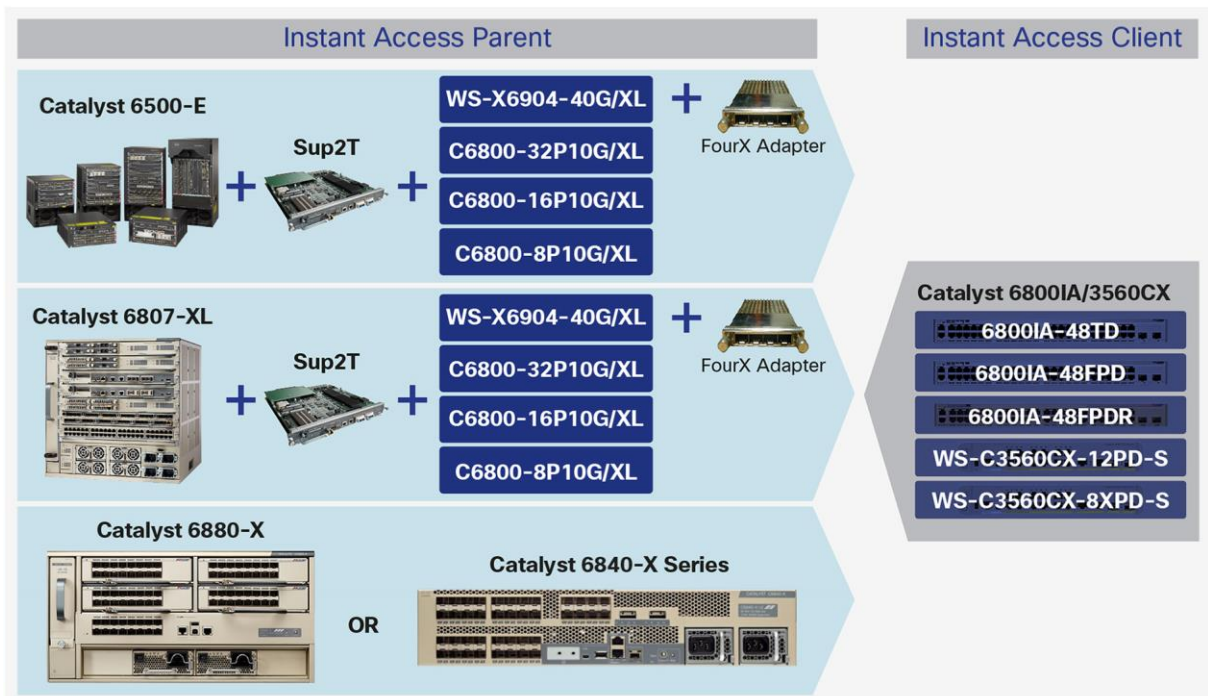
## Cisco Catalyst Instant Access Architecture

**Control Plane**

The control plane implementation in the Instant Access solution allows for the logical grouping of all access switches into one entity. The control plane has four main components:

- **Satellite Discovery Protocol (SDP).** This link-based protocol runs on every link between the Instant Access parent and clients. It establishes, monitors, and maintains fabric link connectivity and allows for a Multichassis EtherChannel connection across parent and client. SDP configures fabric uplinks at the client with no human intervention, providing zero-touch client installation.

- **Satellite Registration Protocol (SRP).** This protocol registers the Instant Access client and performs an image check and automatic upgrade of the client to match the image on the Instant Access parent. This occurs for both new clients and new client stack members as they are added to the stack. SRP provides the ability for online insertion and removal (OIR) and auto provisioning of the client. SRP removes the need for image management at the access layer, which provides the added benefit of Cisco IOS Software feature consistency across the distribution and access layers.

- **Satellite Configuration Protocol (SCP).** This protocol handles configuration management, metrics, and status of Instant Access clients.

- **InterCard Communications (ICC).** ICC is used for infrastructure features like Syslog, QoS, remote login and PoE+ across the Instant Access parent and client.

These control protocols run transparently and automatically in the background. No additional user configuration is required.

## VNTAG

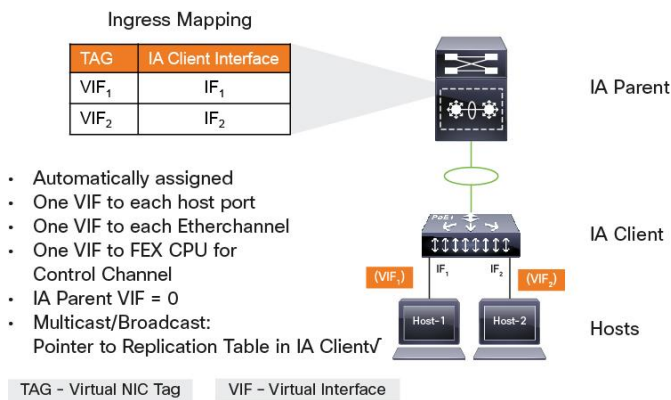A 6-byte VNTAG header is encapsulated on every frame that traverses the fabric link between the Instant Access client and parent as shown in Figure 7. The VNTAG header enables the Instant Access client to behave like a remote line card, allowing client host ports to appear as logical interfaces at the parent switch. To differentiate between unicast packet and multicast packet, the P-bit is used.

**Figure 7.** VNTAG Header
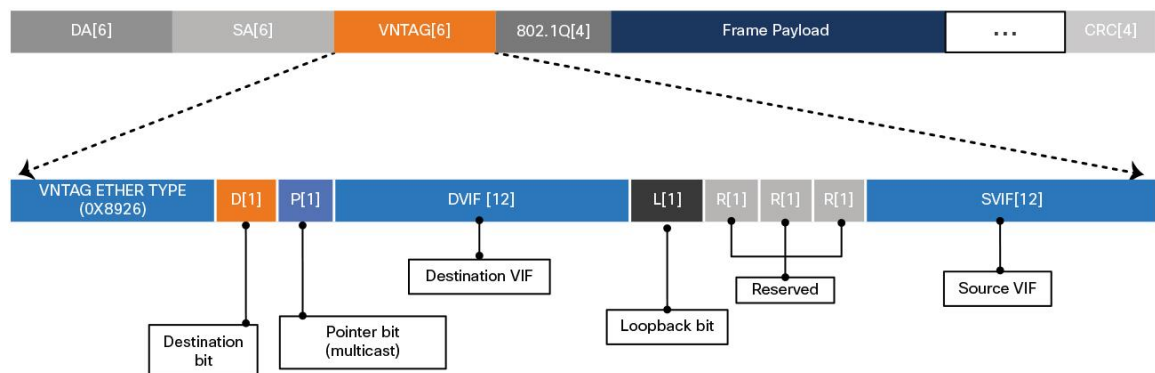
For an Instant Access client to operate as a remote line card to the parent, SRP associates each host port on the client with a unique virtual interface ID (VIF). The Instant Access parent assigns a VIF to each host port on the client during the provisioning process (Figure 8). Any packet that enters the client access switch is tagged with a VNTAG header before being sent to the parent over the fabric links. The VIF assigned to the ingress port is used as the source VIF in this VNTAG header. Conversely, for packets destined for a client switch, the parent uses the destination VIF in the VNTAG header to define the egress port on the client.

**Figure 8.**   VIF Assignment



Ingress Mapping

| TAG | IA Client Interface |
| --- | --- |
| $VIF_1$ | $IF_1$ |
| $VIF_2$ | $IF_2$ |

- Automatically assigned
- One VIF to each host port
- One VIF to each Etherchannel
- One VIF to FEX CPU for Control Channel
- IA Parent VIF = 0
- Multicast/Broadcast: Pointer to Replication Table in IA Client√

IA Parent

IA Client

Hosts

TAG – Virtual NIC Tag        VIF – Virtual Interface

## Unicast Forwarding

To understand unicast traffic flow in the Cisco Catalyst Instant Access solution, following is an example of a unicast packet walk (Figure 9).

1.  A regular Ethernet frame arrives at the Instant Access client host port. For this example, we will refer to this host port as $IF_1$ having VIF = $VIF_1$.

2.  The ingress Ethernet frame is encapsulated with a VNTAG header with source VIF = $VIF_1$ and destination VIF = 0. (All packets that enter at the Instant Access client host port are sent upstream to the Instant Access parent with destination VIF =0.)

3.  A packet with a VNTAG header arriving at the FEX interface at the Instant Access parent is de-encapsulated of the header. The MAC learning happens at the IA parent post VNTAG de-encapsulation. The original Ethernet frame is then processed by the forwarding engine of the parent Catalyst switch and switched like a regular Ethernet frame arriving on a native port.

**Figure 9.**  VNTAG Packet Across Fex-Fabric from Client to Parent



VNTAG - Virtual NIC Tag    VIF - Virtual Interface

4. For packets coming from the core layer toward the Instant Access client host port $VIF_1$, the Instant Access parent does the table lookup (Figure 10). It identifies the outbound fabric link interface to be a FEX, encapsulates the frame with VNTAG header with source VIF = 0 and destination VIF = $VIF_1$, and sends it over the fex-fabric.

5. The Ethernet frame arriving on the fex-fabric at the Instant Access client is de-encapsulated of its VNTAG header, and based on the destination $VIF_1$, is switched to the corresponding interface $IF_1$.

**Figure 10.**  VNTAG Packet Across Fex-Fabric from Parent to Client Host Port



VNTAG - Virtual NIC Tag    VIF - Virtual Interface

This enables the simplicity of the Instant Access solution: The VNTAG is local to the fex-fabric link between the client and parent, and the rest of the network is unaware of it.

## Multicast Forwarding

The Cisco Catalyst 6800ia Series Switch includes intelligent multicast capabilities of local multicast replication in addition to all the multicast capabilities of the Cisco Catalyst 6500 and 6800 Series switches, such as Label Switched Multicast or Medianet. Figure 11 shows how Instant Access performs local multicast replication when multiple receivers are joined at Instant Access client host ports.

1. Multicast group receivers connected to Instant Access client interfaces $IF_1$ and $IF_2$ join multicast groups as part of (*, G)/(S, G) entries at the Instant Access parent.

2. The Instant Access parent programs the Instant Access client for the group VIF table, which maintains the mapping of multicast group VIFs specific to the client's group receiver interfaces.

3. A single copy of each multicast packet is sent over the fex-fabric toward the Instant Access client with destination VIF = group VIF of the multicast group and P bit set to 1 indicating it is a multicast packet.

4. The Instant Access client receives the VNTAG-encapsulated packet with destination VIF = group VIF and the "P" bit set in the VNTAG header to indicate it is a multicast packet. The Instant Access client looks up the group VIF to interface ID mapping table and performs local replication of the multicast packets before sending a copy of each packet to each interface ($IF_1$ and $IF_2$) connected to the receiver.

**Figure 11.** Multicast Packet Replication at the Instant Access Client
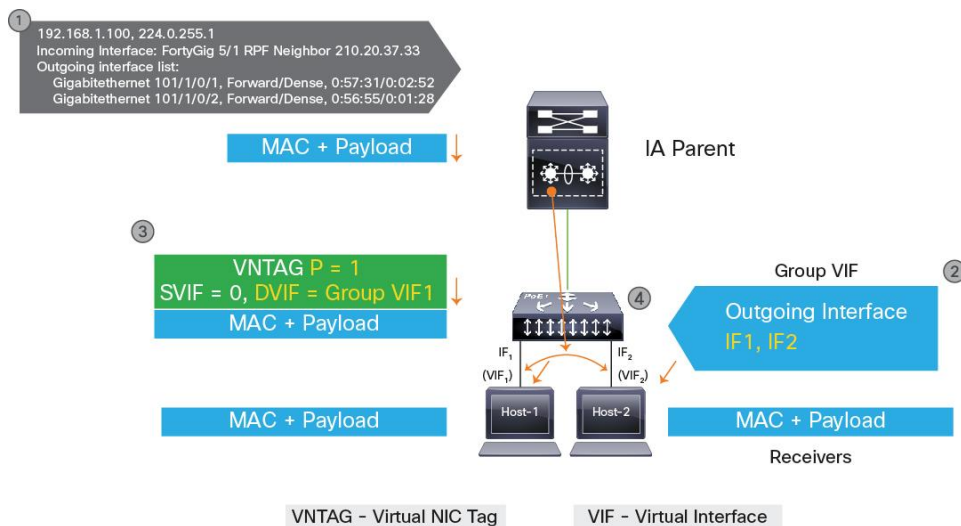


## Solution Capabilities

**Operational simplicity:** The Instant Access solution provides a single point of management across distribution and access switches. All the access host interfaces are represented logically at the Instant Access parent in a four-level interface (Figure 12).

**Figure 12.** Interface Naming



For example, as indicated in Figure 13, a Gigabit Ethernet interface on a Catalyst 6800ia configured as FEX 111 stack member 2 is logically represented at the Instant Access parent as:

```
interface GigabitEthernet111/2/0/1
```

**Figure 13.**    Interface Numbering



interface GigabitEthernet
111/1/0/1

interface GigabitEthernet
111/2/0/1

Every physical host port on each Instant Access client is a logical interface that can be configured and managed locally at the Instant Access parent as shown in the following interface output.

```
Cat6500-VSS#show int gig 111/2/0/1
GigabitEthernet111/2/0/1 is up, line protocol is up (connected)
   Hardware is C6k 1000Mb 802.3, address is 0000.0000.0001 (bia 0000.0000.0001)
  MTU 1500 bytes, BW 1000000 Kbit, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input never, output never, output hang never
  Last clearing of "show interface" counters 3w4d
  Input queue: 0/2000/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     0 packets input, 0 bytes, 0 no buffer
     Received 0 broadcasts (0 multicasts)
     0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
     0 watchdog, 0 multicast, 0 pause input
     0 input packets with dribble condition detected
     0 packets output, 0 bytes, 0 underruns
     0 output errors, 0 collisions, 0 interface resets
```

Similarly, the complete configuration of an Instant Access client FEX ID 111 stack member 2 is logically centralized at the Instant Access parent, as shown in the following output.

```
Cat6500-VSS#show run fex 111 module 2
Building configuration...

Current configuration : 5554 bytes
```

```
!
interface GigabitEthernet111/2/0/1
 switchport access vlan 90
 switchport voice vlan 91
 switchport host
!
interface GigabitEthernet111/2/0/2
 switchport access vlan 90
 switchport voice vlan 91
 switchport host
!
…
```

Since each Instant Access client is treated like a line card to the Instant Access parent, it renders like a line card in the "show module" output.

```
Cat6500-VSS#show module fex 111
 Switch Number:  111   Role:                   FEX
 --------------------  ----------------------------
Mod Ports Card Type                              Model              Serial No.
--- ----- -------------------------------------- ------------------ -----------
  1   48  C6848ia 48GFPwr 2SFP                   C6800IA-48FPD      FHH1707P00S
  2   48  C6848ia 48GFPwr 2SFP                   C6800IA-48FPD      FHH1707P010

Mod MAC addresses                       Hw     Fw           Sw           Status
--- ---------------------------------- ------ ------------ ------------ -------
  1  0022.bdf4.6600 to 0022.bdf4.6633   7.0   15.0(2.0.57) Ok(FLIC Enabled)
  2  0022.bdf4.6d80 to 0022.bdf4.6db3   7.0   15.0(2.0.57) Ok(FLIC Enabled)

Mod  Online Diag Status
---- ------------------
  1  Pass
  2  Pass
```

With up to 42 48-port switches managed like a remote line card and all 2000 ports represented logically at the Instant Access parent, it enables a single point of configuration and management for the entire distribution block.

An example of a stack of five 6800ia switches is shown below:

```
Switch#show module fex 101
 Switch Number:  101   Role:                   FEX
 --------------------  ------------------------------------------
Mod Ports Card Type                              Model            Serial No.
--- ----- -------------------------------------- ------------------
  1   48  C6800IA 48GE                           C6800IA-48TD     FOC1737W0PF
  2   48  C6800IA 48GE POE                       C6800IA-48FPD    FOC1736Z036
  3   48  C6800IA 48GE                           C6800IA-48TD     FOC1737W0NP
  4   48  C6800IA 48GE POE                       C6800IA-48FPD    FOC1741S58N
```

```
    5   48  C6800IA 48GE POE                          C6800IA-48FPD    FOC1736Z03L
  <snip>
```

In addition to stacking clients, it is possible to connect both types of client switches (6800ia and 3560CX) to the same parent switch using different FEX IDs. An example of mixed clients in Instant Access parent switch is shown below:

```
6880X-VSS#show fex
FEX         FEX             FEX                     FEX
Number      Description     State       Model                   Serial
-------------------------------------------------------------------------
105         FEX0105         online      C6800IA-48FPD    FOC1741Y004
107         FEX0107         online      C6800IA-48FPD    FCW1827B0FC
199         FEX0199         online      WS-C3560CX-12PD-S FOC1839Z10H
6880X-VSS#
```

## Configuring Compact Switch in Standalone or Instant Access mode

The 6800ia client works only in Instant Access mode and is provisioned and configured from the parent switch. The 3560CX client can work both as a standalone switch and as an Instant Access client. The mode of operation is configurable, and requires a reload of the client switch.

The mode of operation can be configured directly from the console of the 3560CX switch, using the "fex-mode enable" and "fex-mode disable" commands. It can also be configured directly from the parent Catalyst 6500/6800 parent switch when connected. The command "show fex-mode" on the client switch displays the current mode on the switch.

The following example shows the conversion procedure on the Cisco Catalyst 3560-CX Series switch.

```
3560CX-12PD#show fex-mode
Switch is in non Fex mode
3560CX-12PD#
3560CX-12PD#fex-mode ?
  disable  Disable Fex mode
  enable   Enable Fex mode

3560CX-12PD#fex-mode enable
System will reload after mode conversion.
Do you want to continue? [no]: yes
3560CX-12PD#
```

After reload, the switch comes up as an Instant Access client and can be provisioned and managed from the parent switch similar to the 6800ia client.

Similarly, to change the mode back to standalone, the "fex-mode disable" command can be used, which also requires a reload before the switch comes up in standalone mode.

When the 3560CX switch is connected to the parent with Instant Access, it can be converted to standalone mode directly from the parent switch. The command "reload fex <fex-id> standalone" can be used to convert a specific client into standalone mode, or "reload fex all standalone" can be used to change the mode on all clients capable of operating in standalone mode.

```
6880X-VSS#show fex
FEX         FEX             FEX                     FEX
Number      Description     State           Model               Serial
-------------------------------------------------------------------
105         FEX0105         online          C6800IA-48FPD       FOC1741Y004
107         FEX0107         online          C6800IA-48FPD       FCW1827B0FC
199         FEX0199         online          WS-C3560CX-12PD-S   FOC1839Z10H
6880X-VSS#
6880X-VSS#reload fex 105 standalone
FEX 105, module 1 doesn't support Standalone conversion
6880X-VSS#
6880X-VSS#reload fex 199 standalone
Proceed with reload of fex module and Convert to Standalone mode?[confirm]
6880X-VSS#
```

Another method to convert the client into standalone mode from the parent switch is to attach to the client using the "attach fex-id <fex-id>" command and then configuring "fex-mode disable". The status of the client when attached to it in FEX mode is shown below:

```
FEX-199#show fex-mode
FEX             FEX             FEX             FEX
Number      Description     State           Model               Serial
----------------------------------------------------------------------------
---
199         Local FEX       online          WS-C3560CX-12PD-S   FOC1839Z10H
```

Release 15.2(1)SY1 adds an additional parameter to the process to convert the compact switch to standalone mode. It allows the user to also disassociate the client switch from Instant Access while being converted into standalone mode. This can be configured with the command "reload fex <fex-id> standalone dissociate".

## Instant Access with Multigigabit Support

With the new client, WS-C3560CX-8XPD-S, there are two interfaces available on the Instant Access client switch that support Multigigabit (mGig) Ethernet speeds:

```
DIST-VSS#sh int status fex 103
Port            Name        Status      Vlan        Duplex  Speed Type
Gi103/1/0/1                 disabled    1           full    auto 10/100/1000BaseT
...
Gi103/1/0/6                 disabled    1           full    auto 10/100/1000BaseT
Te103/1/0/7                 disabled    1           a-full  auto
100/1G/2.5G/5G/10GBaseT
Te103/1/0/8                 disabled    1           a-full  auto
100/1G/2.5G/5G/10GBaseT
DIST-VSS#
```

```
DIST-VSS(config)#int t103/1/0/8
DIST-VSS(config-if)#speed ?
  100    Force 100 Mbps operation
  1000   Force 1000 Mbps operation
  10000  Force 10000 Mbps operation
  2500   Force 2500 Mbps operation
  5000   Force 5000 Mbps operation
  auto   Enable AUTO speed configuration
DIST-VSS#show int t103/1/0/8 status
Port            Name    Status      Vlan      Duplex  Speed Type
Te103/1/0/8             connected   1           full  a-10G
100/1G/2.5G/5G/10GBaseT
DIST-VSS#
```

## Automatic Provisioning of Access Clients

Instant access further simplifies the initial provisioning of the access layer by automatically provisioning the Instant Access clients as they connect to the fabric links of the Instant Access parent. The Instant Access parent discovers the Instant Access client and also performs the software image upgrade if the client image is not the same as that of the Instant Access parent. Both of these actions occur automatically, without any user intervention. The Instant Access client uses the FlexStacking-Plus stacking protocol to enable stacking between members with 80 Gbps of bidirectional stack bandwidth and up to five Instant Access clients in a stack. Just like FlexStack-Plus, the stack master is automatically elected and new stack members are discovered and provisioned automatically by the Instant Access parent, truly like a line card to the parent switch.

Pre-provisioning the Instant Access client switch configuration before physical installation is supported. Once an Instant Access client is connected, the pre-provisioned configurations are applied to the Instant Access client host ports automatically, further simplifying deployment: A network administrator can pre-provision Instant access clients from the network distribution layer and have the Instant access clients installed and cabled by anyone locally who does not need to be networking-savvy.

Following is an example where an Instant Access client (FEX 112) is pre-provisioned as a stack of two.

```
module provision  create fex  112 type 6800IA-48TD
module provision  create fex  112 module 2 type 6800IA-48TD

Config# Interface range 112/1/0/1 – 3
Config# switchport access vlan 100
Config# switchport voice vlan 101
Config# switchport host
```

Once the client ID (FEX-ID 112) is pre-provisioned, the Instant Access client configuration for interface host ports shows up in the running configuration at the Instant Access parent. This configuration can be checked by issuing the command "show run fex 112".

```
Cat6500-VSS#show run fex 112
Building configuration...


Current configuration : 11103 bytes
```

```
!
interface GigabitEthernet112/1/0/1
 switchport access vlan 100
 switchport voice vlan 101
 swtichport host
!
interface GigabitEthernet112/1/0/2
 switchport access vlan 100
 switchport voice vlan 101
 switchport host
!
interface GigabitEthernet112/1/0/3
 switchport access vlan 100
 switchport voice vlan 101
 switchport host
!
```

As the new Instant Access client is physically connected, the control protocols automatically configure the client uplinks to the parent, and then the pre-provisioned configuration is automatically applied to the client's host port interfaces.

## Scalability with Instant Access

The Instant Access solution is designed to support an optimal number of host ports from a single point of management. The solution needs to be able to effectively scale to manage a typical deployment, while at the same time not overwhelming system resources on the parent switch, to provide a stable and efficient system. The total number of host ports supported with Instant Access is determined by the control plane or CPU resources of the parent switch, as well as the capabilities on the client switch for parameters such as stacking. Table 1 summarizes the support for different Instant Access systems.

**Table 1.**     Scalability with Instant Access

| Scalability Parameter | 6500-E/6807-XL (Supervisor 2T) with 6800ia 15.1SY Train | 6880-X with 6800ia 15.2SY Train | 6500-E/6807-XL (Supervisor 2T) with 6800ia 15.2SY Train | 6500-E/6807-XL or 6880-X with 3560-CX 15.2SY Train |
|---|---|---|---|---|
| **Maximum host ports** | 1008 | 2016* | 1200 with 15.2(1)SY 1536 with 15.2(1)SY1 | 504 (12-port switch) 336 (8-port switch) with 15.2(1)SY1 |
| **Maximum FEX IDs** | 12 | 42 | 25 or 32 | 42 |
| **Maximum client switches** | 21 | 42 | 25 or 32 | 42 |
| **Maximum clients in stack** | 3 | 5 | 5 | N/A |
| **Maximum user ports in stack** | 144 | 240 | 240 | N/A |
| **Maximum bandwidth of fabric link** | 60 Gbps | 80 Gbps | 80 Gbps | 20 Gbps |

**Note:**   Catalyst C6840-X family of switches support a total of 1536 host ports when used as a parent switch with Instant Access.

As noted in the table, the scalability of the Instant Access solution increased with the 15.2(1)SY software version. The Cisco Catalyst 6880-X Switch and Supervisor Engine 2T offer different levels of scalability due to their differing CPU capabilities. The 12-port 3560-CX switch, when used as an Instant Access client (with a 6880-X or Supervisor 2T based parent switch), supports a maximum of 504 ports (42 client switches, with 12 ports each).

The term "FEX ID" denotes an Instant Access client stack or a Fabric PortChannel. With earlier releases, the 21 clients had to be deployed in such a way that there were a total of 12 FEX IDs or 12 stacks. With the 15.2(1)SY release, there is no restriction on how the 42 clients can be deployed. (They can all be standalone clients or configured in stacks.)

The following output shows a fully scaled Instant Access system and the corresponding platform resources used.

```
6880X-VSS#show fex system platform usage
FEX id usage details
<snip>


FEX slot usage details
   FEX-id    Switch-id    Vslot    Pslot    Status
   ------    --------     -----    -----    ------
    101        23          91        1      In-use
    <snip>
    121        22          87        5      In-use
   Total      Used      Reserved    Free
   -----      ----      --------    ----
   42          42          0          0


FEX ports usage details
   FEX-id    Switch-id    Ports
   ------    --------     -----
    107        3           48
    <snip>
    101        23          48
   Total      Used       Free
   -----      ----       ----
   2016       2016         0


Stack members usage details
   FEX-id    Switch-id    Used    Free
   ------    --------     ----    ----
    107        3           1       4
    106        4           3       2
   <snip>



VNTAG MGR Usage
----------------------
```

```
Max unicast VIFs available      2048
Total unicast VIFs used         2016
Max non-mdest VIFs available    1019
Total non-mdest VIFs used       59
Max mdest VIFs available        16380
Total mdest VIFs used           2409


LTL MGR Usage
------------------
MAX unicast LTLs available  2048
Total unicast LTLs used     2016
6880X-VSS#
```

## Simplified Software Management

The Cisco Catalyst 6500/6800 software image and Instant Access client image are bundled as a single image, truly like a line card image at the parent. Whenever a new Instant Access client boots up and is discovered by the Instant Access parent, it automatically checks if the Instant Access client image matches the software image on the Instant Access parent. If it does not match, the Instant Access parent updates the client image automatically. This eliminates the need to perform software upgrades at the access layer and enables an agile infrastructure with consistent features across distribution and access layers. The single image also includes the images for all the client types supported (6800ia and 3560CX platforms), facilitating initial deployments, upgrades and replacements.

## High Availability

The Instant Access solution provides multiple levels of resiliency. At the distribution layer, the Instant Access parent supports the Cisco Virtual Switching System (VSS) and Cisco Virtual Switching System Quad-Supervisor (VS40) configurations, providing high availability from any point of failure. With Quad-supervisor SSO at the distribution level, it would take three supervisor failures before losing network connectivity when Instant Access client stacks are dual-homed to the Instant Access parent (Figure 14).

**Figure 14.**   Instant Access Parent High Availability
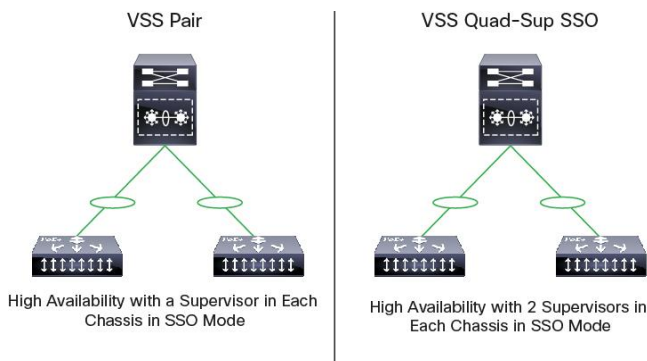


The multiple fabric links bundled into a Multichassis EtherChannel connection between parent and client can scale up to 80 Gbps with eight 10 Gigabit Ethernet links between the VSS pair and the client stack, providing fabric link redundancy. The fabric link can span across stack members, providing redundancy as well. The Instant Access parent and client support EtherChannel load sharing over the fex-fabric to provide a high level of redundancy across multiple fex-fabric links (Figure 15).

**Figure 15.** Fex-Fabric High Availability



Dual Homed to
VSS Pair

Dual Homed across
Stack Members

Up to 8 Uplinks (80G)
MEC across Client to Parent

The initial release of Instant Access supported stacking of up to three switches, enabling six 10 Gigabit Ethernet links between the VSS pair and client stack. With the increase in stacking support to five switches, up to 10 uplinks are available to be used. Because a maximum of eight interfaces can be bundled in an EtherChannel, any eight of the ten 10 Gigabit Ethernet uplinks can be used to provide an 80 Gbps fabric connection:

```
Switch#show etherchannel 10 summary
 Flags:  D - down       P - bundled in port-channel
 !
 Number of channel-groups in use: 3
 Number of aggregators:          3
 Group  Port-channel  Protocol    Ports
 ------+------------+------------------------------
 10     Po10(SU)        -         Te1/2/5(P)    Te1/2/6(P)    Te1/2/7(P)
                                  Te1/2/8(P)    Te2/2/5(P)    Te2/2/6(P)
                                    Te2/2/7(P)     Te2/2/8(P)


 Last applied Hash Distribution Algorithm: Adaptive


Switch#show fex 101 detail
FEX: 101        Description: FEX0101     state: online
FEX version: 15.2(3.2.3)E
Extender Model: C6800IA-48FPD, Extender Serial: FOC1736Z036
FCP ready: yes
Image Version Check: enforced
Fabric Portchannel Ports: 8
Fabric port for control traffic: Te1/2/5
Fabric interface state:
Po10      - Interface Up.
    Te1/2/5   - Interface Up.      state: bound
    Te1/2/6   - Interface Up.      state: bound
    Te1/2/7   - Interface Up.      state: bound
    Te1/2/8   - Interface Up.      state: bound
    Te2/2/5   - Interface Up.      state: bound
```

```
Te2/2/6    - Interface Up.      state: bound
Te2/2/7    - Interface Up.      state: bound
Te2/2/8    - Interface Up.      state: bound
```

The Instant Access client supports host-port EtherChannels downstream from the client switch. Up to two Instant Access client ports can be members of a host-port EtherChannel. While it is possible to configure up to eight member interfaces in a host-port EtherChannel, earlier software releases support only a two-member EtherChannel. With 15.2(1)SY1, an eight-member host-port EtherChannel is also supported with Instant Access. The EtherChannel can span across stack members in an Instant Access client, but not across different Instant Access client stacks (Figure 16). This functionality is supported on both the Cisco Catalyst 6800ia and the 3560-CX Series clients. A total of 23 host-port EtherChannels are supported on each Instant Access client.

**Figure 16.**  Host-Port High Availability



MEC at IA Host Ports
Across Stack Members

## Enhanced Fast Software Upgrade

Cisco Catalyst 6500E and 6800 Series switches support enhanced Fast Software Upgrade (eFSU). This increases network availability by reducing the downtime caused by software upgrades across two supervisor engines in a VSS pair. The upgrade brings the active and standby supervisors into synchronous Stateful Switchover (SSO) mode across two supervisor engines running two different software versions. It maintains an active data plane on both switches in the VSS pair, providing increased network availability during the upgrade process.

eFSU is a four-step process (Figure 17):

**Step 1.** issu loadversion command: The new software image is loaded on the standby supervisor on the VSS pair.

**Step 2.** issu runversion command: The new software is loaded on the standby supervisor engine while the active supervisor engine continues to operate with the previous software version. As part of the upgrade, the standby supervisor reaches the SSO hot-standby stage, a switchover occurs, and the standby becomes active, running the new software version.

**Step 3.** You can continue with the upgrade to load the new software on the other processor with the issu acceptversion command, or you can abort the upgrade and resume operation with the old software with the issu abortversion command.

**Step 4.** issu commitversion command: This command completes the process of eFSU by loading the new software version on the standby supervisor engine. For more details about eFSU, click here.

**Figure 17.** eFSU Steps



eFSU capability is extended to support Instant Access client upgrades similar to how a line card is upgraded. The client software image is bundled with the Catalyst 6500 or 6800 Series software image. A new command-line interface (CLI) is introduced, enabling the upgrade of the Instant Access client stack (FEX IDs), which in turn enables an upgrade of the Instant Access client's software version before the issu commitversion command (after step 3 and before step 4) of the eFSU process.

```
issu runversion [fex[range] <num | all >]
```

The issu runversion fex command initiates the upgrade of the Instant Access client's move to a new software version. A user can specify a set (or range) of FEX IDs for the rolling upgrade and a reload of Instant Access clients. After all clients are upgraded, a user has the choice to abort the eFSU process and go back to the previous software version using or completing the eFSU process with the issu commitversion command (see Figure 18).

**Figure 18.** eFSU Instant Access Client Upgrade
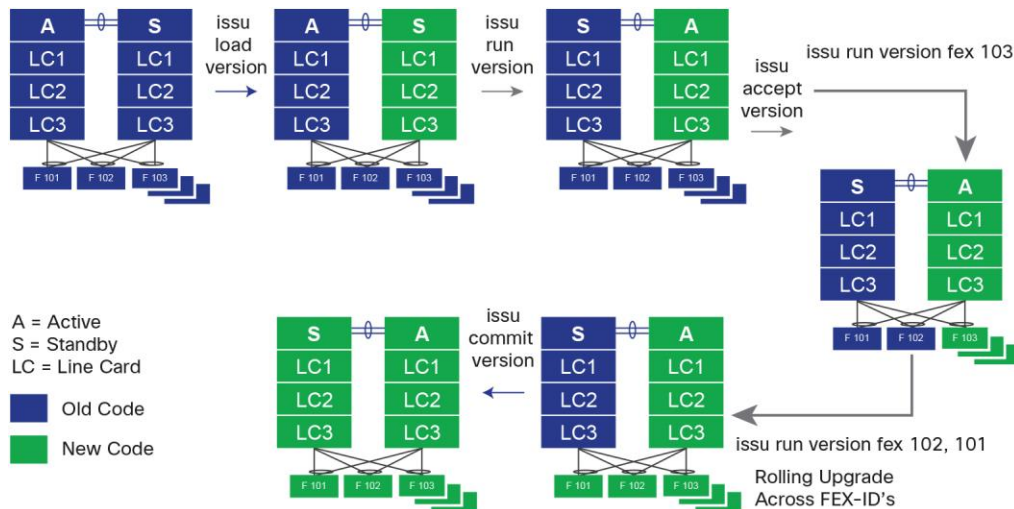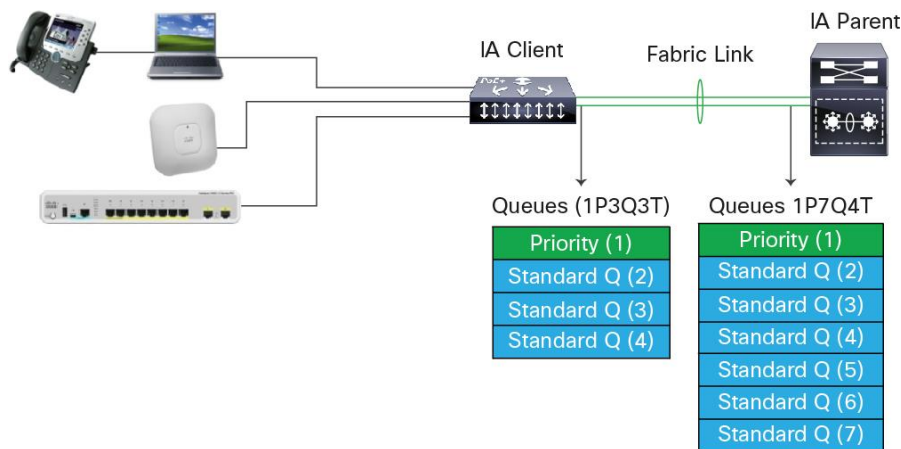


## Quality of Service

With a stack of three switches, the Instant Access solution provides up to 60 Gbps of fex-fabric uplink connectivity per stack (of three Instant Access clients) to the VSS pair, offering the subscription ratio of 2.4 to 1. When the stacking capability is increased to five switches, and the fex-fabric uplink connectivity per stack increases to 80 Gbps, the subscription ratio is 3 to 1.

Instant Access client fabric links support four queues (1P3Q3T), with one priority queue and three standard queues. The line card on the Instant Access parent supports eight (1P7Q4T) queues on the fabric link (Figure 19).
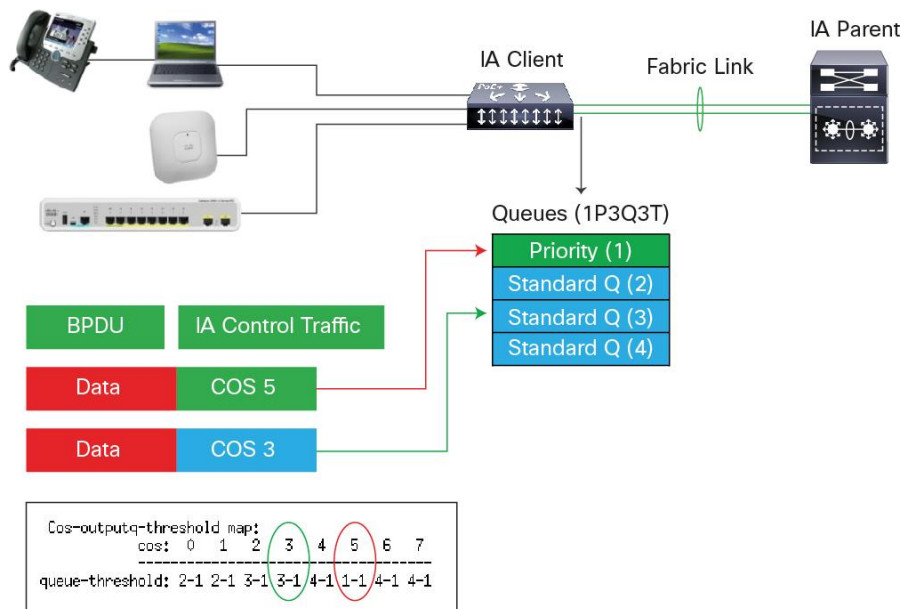
**Figure 19.** QoS Queues at the Instant Access Client and Parent



IA Parent

IA Client   Fabric Link

Queues (1P3Q3T)        Queues 1P7Q4T

| Priority (1) |
| Standard Q (2) |
| Standard Q (3) |
| Standard Q (4) |

| Priority (1) |
| Standard Q (2) |
| Standard Q (3) |
| Standard Q (4) |
| Standard Q (5) |
| Standard Q (6) |
| Standard Q (7) |

Quality of service (QoS) over the fabric link is strictly based on differentiated-services-code-point (DSCP)/class-of-service (CoS) values of the ingress packets. The Instant Access parent and client maintain a default DSCP-to-Queue map and CoS-to-Queue map, which is the basis of queuing packets appropriately over priority queue or standard queue on fex-fabric interfaces.

As Figure 20 shows, any IP packet marked with COS=5 is queued over priority queue 1, and any IP packet marked with COS=3 is queued to standard queue 3. All Instant Access control traffic is also sent over the priority queue to ensure that communication between the Instant Access parent and client is not lost due to congestion.

**Figure 20.** QoS at Instant Access Client



IA Parent

IA Client   Fabric Link

Queues (1P3Q3T)

| Priority (1) |
| Standard Q (2) |
| Standard Q (3) |
| Standard Q (4) |

| BPDU | IA Control Traffic |
| Data | COS 5 |
| Data | COS 3 |

```
Cos-outputq-threshold map:
          cos:  0   1   2   3   4   5   6   7
              --------------------------------
queue-threshold: 2-1 2-1 3-1 3-1 4-1 1-1 4-1 4-1
```

After IP packets arrive over the fabric link at the Instant Access parent, they can be marked, remarked, classified, or policed. Likewise, traffic downstream from the Instant Access parent over the fabric port that is heading toward the Instant Access client host port uses default DSCP-CoS to queue maps and traffic in the appropriate queue. Cisco IOS Software Release 15.2(1)SY1 introduces new QoS functionality on the Instant Access host ports.

The following parameters are now configurable on client host ports:

- Priority Queuing
- Queue Bandwidth
- DSCP to queue map
- Queue limit/Buffer

They can be configured using Modular QoS CLI, and the QoS configuration applies to the entire Instant Access client stack. More details can be found here.

## Consolidated Security Features

When building a campus network, the number one issue that comes to mind is usually security. Cisco Instant Access supports Cisco TrustSec®, inheriting the Catalyst 6500/6800 capabilities.

Instant Access provides a single consistent security policy across the enterprise campus network. The solution supports:

- Role-based access control with Security Group Tagging (SGT)
- Security Group Access List (SGACL)
- IP subnet, VLAN, and port-based SGT mapping
- Network Device Admission Control (NDAC)
- 802.1x, WebAuth, and Mac Authentication Bypass (MAB) authentication for identity
- IBNS 2.0 framework of features, including Common Classification Policy Language (C3PL)-based configuration

All the security policies are applied at the IA parent only with no configuration at the access layer. Access lists are enforced at IA parent only. Any packet arriving at IA client host port are VNTag-ed and sent to IA parent which decapsulates the VNTag and enforces the access list policies on it (Figure 21).

**Figure 21.** Inbound Access List
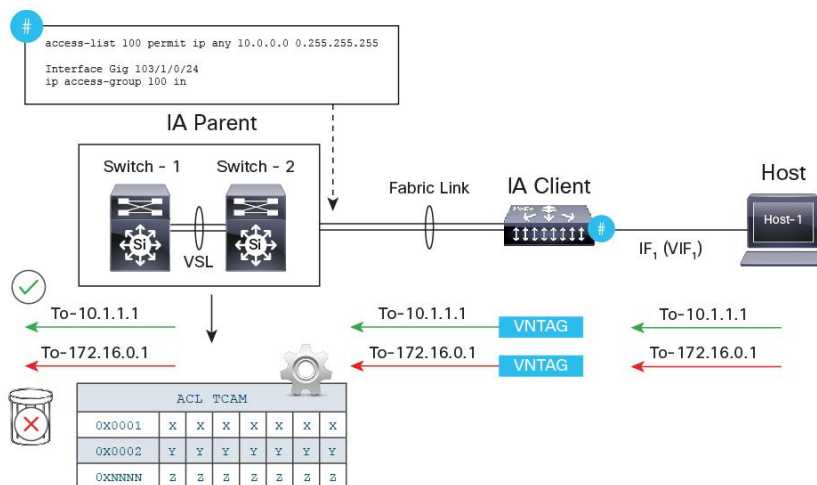


Similarly, for packets arriving at IA parent and egressing the IA client host ports, the policies are applied at IA parent before the packet is switched over fex-fabric link to IA client (Figure 22).

**Figure 22.** Egress Access List
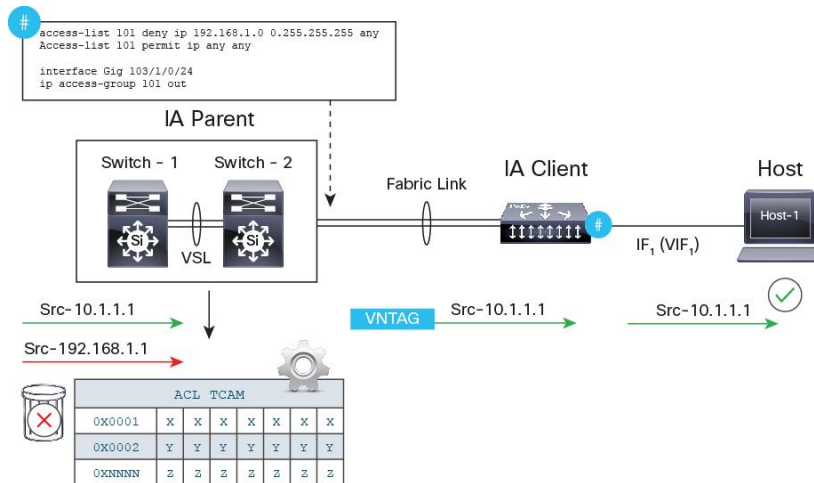


The Instant Access parent acts as both the Security Group Tag (SGT) imposition point and Security Group Access List (SGACL) enforcement point (Figure 23). Cisco ISE communicates with the Instant Access parent and enforces policies that are configured by the network administrator in the Cisco ISE. Instant Access also supports SGT & SGACL based policies based on IP subnet, VLAN, or a Layer 3 port in absence of Cisco ISE in the network.

Instant Access supports Network Device authentication (NDAC) guaranteeing the physical infrastructure is secure. Network device authentication is done at IA parent only and is not required for IA clients, thus reducing the overhead of NDAC authentication at access layer.

The Instant Access client is hardware capable of IEEE MAC Security standard (MACsec), which will be supported in subsequent releases.

Instant Access support 802.1x, MAC authentication bypass, and WebAuth port-based identity services. Instant Access parent communicates with Cisco Identity Services Engine (ISE) controlling the access to the network, and thus enabling single point of management and configuration for all security policies across the network.

**Figure 23.** SGT and SGACL
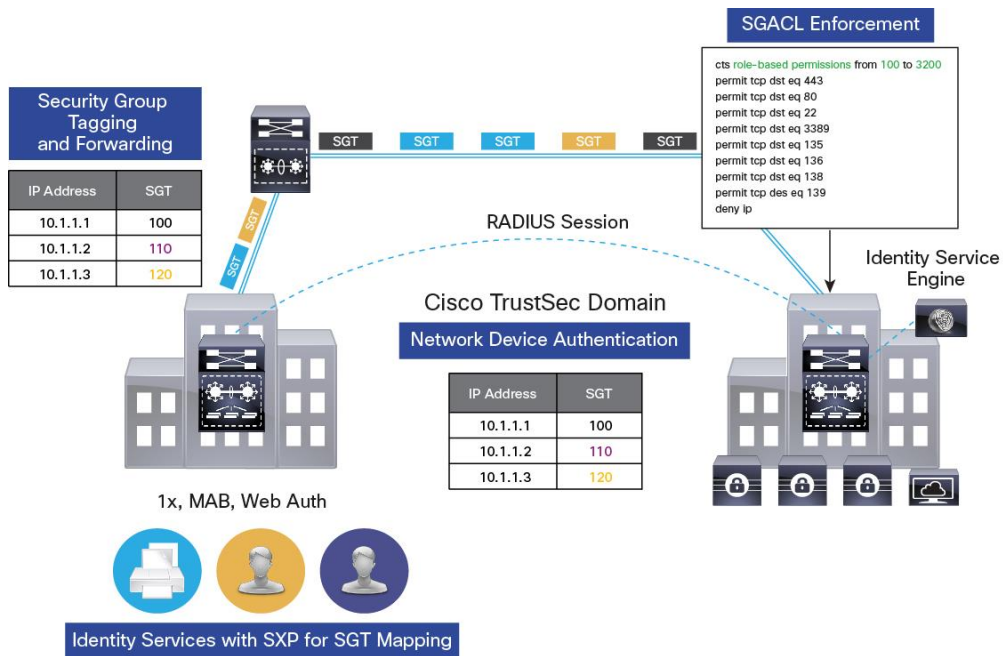


## Unified Application Visibility

Cisco Catalyst Instant Access provides a single point of application visibility and control for a complete distribution block. A single point of configuration and export at the Instant Access parent drastically reduces the complexity of multiple exports from individual access switches and multiple records at the NetFlow Collector (see Figure 24).

**Figure 24.** NetFlow



## Easy VSS, Auto-FEX, and Switch Renumbering

With Release 15.2(1)SY1, there are new features which make the provisioning and management of Instant Access parent and client switches easier.

Easy VSS is a feature by which a VSS system can be configured from two standalone switches in a simplified manner. The feature is not specific to Instant Access, and can be used in any VSS setup. Traditionally, VSS requires that the user configure the following on both switches: assign a Virtual Switch Domain, assign a Switch ID, create a port-channel, configure the port-channel as a virtual switch link (VSL), add interfaces to the VSL port-channel, and lastly, issue "switch convert mode virtual" on both switches. With Easy VSS, we start by connecting the two switches together and making sure that the interfaces that will be a part of the VSL are up and have CDP

(Cisco Discovery Protocol) enabled on them. The following configuration on either of the switches applies the conversion process to convert the pair into a VSS:

To enable (or disable) the feature (on both switches)

```
Switch1(config)#switch virtual easy
To convert to VSS (on any one of the switches)
switch convert mode easy-virtual-switch domain [domain id] links [intf1..intf8]
Switch1#switch convert mode easy-virtual-switch ?
  domain  Select Unique VSL Domain number in your
  Network, Default domain ID is 100
  links    Select VSL Links
```

Auto-FEX

With the Auto-FEX feature in 15.2(1)SY1, the Instant Access parent and client get provisioned automatically upon enabling the interfaces of the parent switch that the client switch is connected on. To enable this feature, the command, " fex auto-config", needs to be configured on the parent switch.

Release 15.2(1)SY1 also supports interface aliases, which can be used in place of the interface ID for commonly used interfaces. The example below illustrates how they can be configured.

Associate an interface with an alias and use this alias name to address the interface:

```
DIST-VSS(config)#interface g101/1/0/48
DIST-VSS(config-if)#alias ?
  LINE  Up to 80 characters describing this interface
DIST-VSS(config-if)#alias blue
DIST-VSS(config-if)#end
DIST-VSS#show interfaces alias all
Interface Name                          Alias
-------------------------------------   --------------------
GigabitEthernet101/1/0/48               blue
DIST-VSS#
DIST-VSS(config)#interface alias blue
DIST-VSS(config-if)#no switchport trunk allowed vlan 1
DIST-VSS(config-if)#end
DIST-VSS#
```

Switch renumbering in a stack is supported with the 15.2(1)SY1 release. For a failed switch that has to be replaced with a return materials authorization (RMA), it is now possible to renumber the stack member while it is connected in the Instant Access client stack. This helps to derive the configuration of the previously failed stack member and apply it to the replacement switch. The example below illustrates this.

```
6880-VSS#module provision update fex 109
6880-VSS(exec-fex-update)#renumber 5 to 4
%FEX 109 slot 5 will reload upon commit.
Are you sure you want to proceed? [no]: yes
6880-VSS(exec-fex-update)#renumber 4 to 5
%FEX 109 slot 4 will reload upon commit.
```

```
Are you sure you want to proceed? [no]: yes
6880-VSS(exec-fex-update)#show
Current module renumber mappings for FEX 109
-------------------------------------------
 renumber 4 to 5
 renumber 5 to 4
Current module Priority mappings for FEX 109
-------------------------------------------
Temp vslots allowed:NO
Current Temp vslot allowed FEXs:
6880-VSS(exec-fex-update)#commit
%FEX 109 renumbered modules will reload.
Are you sure you want to proceed? [no]: yes
6880-VSS(exec-fex-update)#end
6880-VSS#
```

## Interface Templates and AutoConf

The Instant Access solution supports Interface Templates with Cisco IOS Software Release 15.2(1)SY. An interface template is a container of configurations or policies that can be applied to specific interfaces.

All interface templates are customizable and can be easily modified. The template updates immediately ripple to the interfaces and support full rollback functionality. Both per-session and per-port templates are supported, and the solution is compatible with Session Networking or AutoConf features. One of the major advantages of interface templates is that the running configuration will have a fixed and consistent configuration, which in turn reduces the configuration file size.

Interface templates are easy to use, as demonstrated in the following output. They can be statically applied using the source template <template name> command in the CLI. The full interface configuration can be viewed with the show derived-config interface <interface ID> command.

More details on interface templates can be found at this link:
http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/ibns/configuration/15-e/ibns-15-e-book/ibns-int-temp.html

To configure an interface template:
```
DIST-VSS(config)#template IA_TEMPLATE
DIST-VSS(config-template)#switchport mode access
DIST-VSS(config-template)# switchport access vlan 100
DIST-VSS(config-template)# switchport nonegotiate
DIST-VSS(config-template)# switchport port-security
DIST-VSS(config-template)# source template IA_TEMPLATE2
DIST-VSS(config-template)#
DIST-VSS(config-template)#template IA_TEMPLATE2
DIST-VSS(config-template)# spanning-tree portfast edge
DIST-VSS(config-template)#exit
```

To apply an interface template:

```
DIST-VSS(config)#int range g101/1/0/1-12
DIST-VSS(config-if-range)#source template IA_TEMPLATE
DIST-VSS(config-if-range)#end
```

Viewing the derived configuration from an Interface Template:

```
DIST-VSS#show run int g101/1/0/1
Building configuration...
Current configuration : 126 bytes
!
interface GigabitEthernet101/1/0/1
 switchport
 switchport trunk allowed vlan 1
 shutdown
 source template IA_TEMPLATE
end
DIST-VSS#show derived-config int g101/1/0/1
Building configuration...
Derived configuration : 228 bytes
!
interface GigabitEthernet101/1/0/1
 switchport
 switchport access vlan 100
 switchport trunk allowed vlan 1
 switchport mode access
 switchport nonegotiate
 switchport port-security
 shutdown
 spanning-tree portfast edge
end
```

Modifying an Interface Template

```
DIST-VSS(config)#template IA_TEMPLATE
DIST-VSS(config-template)#switchport access vlan 200
DIST-VSS(config-template)#end
DIST-VSS#show derived-config interface g101/1/0/1
Building configuration...
Derived configuration : 228 bytes
!
interface GigabitEthernet101/1/0/1
 switchport
 switchport access vlan 200
 switchport trunk allowed vlan 1
 switchport mode access
 switchport nonegotiate
```

```
 switchport port-security
 shutdown
 spanning-tree portfast edge
end


DIST-VSS#
```

Interface templates can be either built in or user defined, and can be viewed with the following command:

```
DIST-VSS#show template interface brief

Template-Name                               Source          Bound-to-
Interface
------------                                ------          --------------
----
AP_INTERFACE_TEMPLATE                       Built-in        No
DMP_INTERFACE_TEMPLATE                      Built-in        No
IA_TEMPLATE                                 User            Yes
 NESTED TEMPLATE: IA_TEMPLATE2
IA_TEMPLATE2                                User            Yes
IP_CAMERA_INTERFACE_TEMPLATE                Built-in        No
IP_PHONE_INTERFACE_TEMPLATE                 Built-in        No
LAP_INTERFACE_TEMPLATE                      Built-in        No
MSP_CAMERA_INTERFACE_TEMPLATE               Built-in        No
MSP_VC_INTERFACE_TEMPLATE                   Built-in        No
PRINTER_INTERFACE_TEMPLATE                  Built-in        No
ROUTER_INTERFACE_TEMPLATE                   Built-in        No
SWITCH_INTERFACE_TEMPLATE                   Built-in        No
TP_INTERFACE_TEMPLATE                       Built-in        No
DIST-VSS#
```

Templates can be extended to sessions using service templates, which apply to specific access sessions on any given port. A service template contains a set of service-related attributes or features, such as access control lists (ACLs) and VLAN assignments, that can be activated on one, or more, subscriber sessions in response to session events. Both interface templates and service templates can be applied using the AutoConf feature. This involves autoprovisioning of network access based on **who** or **what** is connecting, using identity-based access control or device-based access control. The following output shows the AutoConf policy and built-in parameter map:

```
DIST-VSS#show policy-map type control subscriber BUILTIN_AUTOCONF_POLICY
BUILTIN_AUTOCONF_POLICY
  event identity-update match-all
    10 class always do-until-failure
      10 map attribute-to-service table BUILTIN_DEVICE_TO_TEMPLATE
DIST-VSS#
DIST-VSS#show parameter-map type subscriber attribute-to-service all
Parameter-map name: BUILTIN_DEVICE_TO_TEMPLATE
 Map: 10 map device-type regex "Cisco-IP-Phone"
  Action(s):
```

```
      20 interface-template IP_PHONE_INTERFACE_TEMPLATE
   Map: 20 map device-type regex "Cisco-IP-Camera"
    Action(s):
      20 interface-template IP_CAMERA_INTERFACE_TEMPLATE
   <snip>
   DIST-VSS#
```

The service policy is applied to all AutoConf-enabled interfaces when an identity update event occurs. This event can take place in the form of the detection of a new MAC address, username, user role, device-type classification, or MAC Organizationally Unique Identifier (OUI). The parameter-map BUILTIN_DEVICE_TO_TEMPLATE defines rules against which changes to attributes of the session are evaluated, and an action (such as application of a service template or interface template) is triggered.

More details on AutoConf and configuring it can be found at this link:
http://www.cisco.com/c/en/us/products/collateral/switches/catalyst-3560-x-series-switches/white-paper-c11-732349.html

## Consistent and Rich Features Across the Campus

Table 2 provides a brief list of features that are supported at the Instant Access client host port. For more details on Instant Access and features, click here.

**Table 2.**    Summary of Instant Access Features

| Category | Instant Access |
|---|---|
| Infrastructure | PoE, PoE+, Multichassis EtherChannel, FlexStack |
| Layer 2 | EtherChannel, PAgP, LLDP, (A)VPLS, GRE Tunneling, MPLS, MPLS-VPN |
| IPv6 | IPv6 First- Hop Security, Multicast Routing, QoS, Stateless Auto- Configuration |
| Layer 3 | PBR, EVN, VRF-Lite, PIM SM, WCCPv2, Inter-VLAN Routing, ECMP, Layer 3 Routing Protocols |
| Security | 802.1x Guest VLAN, SXP, SGT, SGACL, IP Source Guard, DHCP Snooping, VACL, RACL, PACL, Flexible NetFlow |
| QoS | Policing, Marking, Rate Limiting, SRR |
| Medianet | Mediatrace, Performance Monitoring |
| Manageability | Autoprovisioning, Interface Templates, AutoConf, Image Management and eFSU |

## Conclusion

Cisco Catalyst Instant Access simplifies the deployment of the enterprise campus network by presenting a single point of configuration, management, troubleshooting, and unified application visibility across the distribution layer. Instant Access also provides consistent features across the campus. The single image management and plug-and-play provisioning of the access layer can enable accelerated rollouts.

## For More Information

For more information, refer to the Cisco Catalyst Instant Access webpage.