

## Cisco IT and the Identity Services Engine

A multiyear deployment journey.

By Greg Rasner  
Security Engagement Manager, Cisco

The Cisco Identity Services Engine (ISE), a policy engine, enables contextual network access control across wired and wireless networks, and extends to mobile connectivity as well (Bring Your Own Device, or BYOD). Contextual controls are based on multiple variables, including who (user identity), when (time of day), where (location), how (access method), and what (device). ISE works with our existing infrastructure to enforce security policy on all devices that attempt to gain access to the network. To do this, ISE can use access switches, wireless controllers, and most Cisco® network gear for edge authentication, as device profiling sensors, and as access enforcement points.

ISE is also capable of extending authentication services on other vendors' 802.1X-compliant hardware, and enabling web authentication as backup for non-802.1X-compliant devices. ISE is deployed as an appliance or runs on a virtual machine (VM). We deploy ISE on a VM, which is in step with our overall data center virtualization and footprint reduction goals. We are taking a measured, controlled approach to rolling out new ISE capabilities. This approach helps IT to ensure a smooth adoption, to collect user feedback, and to build on and leverage ISE capabilities in each phase.

Cisco IT was an early adopter of ISE (deploying ISE 1.1 in 2012), and we have made much progress rolling out ISE capabilities in the last year and a half. See [End-to-End Security Policy Control](#) to learn about our decision-making during the initial deployment phase. That deployment strategy held throughout Cisco's fiscal-year 2014, which ran from August 2013 to July 2014.

This article focuses on key areas of our current ISE deployments, including Deployment Strategy, Testing and Certification Process, Guest Networking and Enhancements, Profiling, Wireless Authentication, Wired Authentication, Replication and Scaling, Operational Support, Pilot and Limited Deployments, and Challenges/Lessons Learned.

## Level Set on Cisco IT's ISE Deployment

At the end of July 2014, we had deployed several primary capabilities for ISE. We have a guest infrastructure managed by two VMs (the primary in our San Jose data center and the secondary in our data center in Western Europe). These VMs manage global guest services for more than 60,000 guest accounts that are created monthly. The guest infrastructure was fully deployed at the end of December 2013 and runs ISE 1.2 Patch 7 (see Table 1). The production cluster runs on 28 VMs and also uses ISE 1.2 Patch 7. On the production cluster, we run all other capabilities. To date we have deployed Wireless Authentication to ISE, 802.1X Monitor Mode and Profiling, at 83 sites globally. We have also deployed Wireless Authentication to ISE on more than 10,000 (of 30,000) Cisco Virtual Office (CVO) systems (these typically run on Cisco 800 Series Routers and Cisco SOHO Routers). And we are running a limited deployment of Wireless Posture Enforcement to two extranet partners for nearly 1000 devices connecting and authenticating successfully. Our profiled endpoint count is currently 550,000 devices, and our maximum concurrent connections at any given time are 55,000.

**Table 1.** IT Deployment Roadmap


Bundle	Capabilities	Scope	ISE version	Deployment Expectation*	Complete
1	ION (Guest Networking)	Global	1.2	Q4FY13	100%
2	Profiling 802.1X wired (Monitor Mode) Endpoint Protection pilot	Global, phased Global, phased 50% ICZ	1.2	Q4FY13	Q3FY15 (Feb 2014)
3	802.1X wireless (Auth Mode)	Global, phased	1.2	Q1FY14	99.99%
4	802.1X wired (Auth Mode)	19 ICZ sites (China)	1.2	Q3FY14	Q1FY15
5	Device onboarding Mobile Device Management (MDM)	Global, phased Global, phased	1.3	Q1FY15	
6	Endpoint Protection SGA Deployment	Global	1.3	Q2FY15	
7	TACACS+ Manager of Managers	Global Global	2.0	Q2FY16 Q2FY16	

\*IT Deployment Dates, not ISE product availability roadmap

## Deployment Strategy

We took a staged approach regarding ISE, to avoid the “Big Bang” of deploying many capabilities simultaneously. To manage capability deployment and adoption, we are implementing ISE capabilities in bundles (see Figure 1). Bundling helps ensure that the requisite steps for success of each capability are completed properly.

**Figure 1.** ISE Capabilities Bundling Strategy



We use a well-defined solution delivery methodology or project lifecycle (PLC) framework that aligns to our PLC. The PLC defines several stages in the delivery of a solution or IT capability, with approval gates and specific deliverables at each stage. All projects follow several phases:

- Business Commit, where the business case and benefits are clearly articulated
- Concept Commit, where the solution architecture is the primary outcome
- Execute Commit, where detailed design work and planning starts

Architecture, Design, and Operational Readiness Reviews occur throughout the process, ensuring that no solution is implemented without the appropriate resources, support, and communications plans in place.

Careful project management, along with governance from the Service Management Office and the Project Management Office, and regular syncing with the project stakeholders ensure that projects stay on time and on budget. Deviations are quickly identified, and appropriate remediation or business decisions are made expediently to get back on track.

Ultimately, the business driver for deploying ISE is our strategy of embedding security within the network fabric. Relying on outdated, hardened, perimeter-based models does not provide the security essential in a borderless enterprise, or in an environment where trends such as BYOD and the Internet of Everything (IoE) result in many more devices that require secure connectivity.

Our global ISE deployment also delivers operational and architectural simplicity, which, in turn, drives real cost savings. Previously, we had to maintain 12 servers globally with long access control lists (ACLs) to provide guest access. With the deployment of ISE Guest, we have reduced our footprint to two VMs globally and an ACL of fewer than 20 lines. Deriving the exact cost savings depends on the total cost of ownership (TCO) calculator. We have seen actual TCO reductions in operational, infrastructure, and upgrade costs for our guest access service alone.

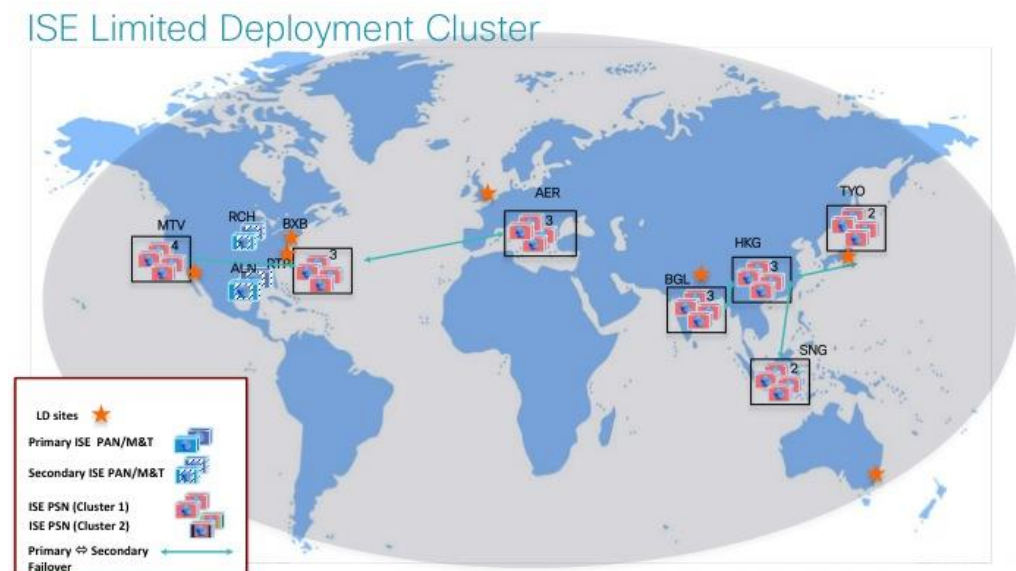
## Testing and Certification Process

Testing and certification for versions and capabilities on ISE has evolved over the last two years. Initially, we used a three-stage process to test and certify releases and patches. First, we use the Service Verification Lab (SVL). Then we test at our San Jose campus, Building 12, which has a separate ISE cluster and Wireless LAN Controllers (WLCs) to ensure that patches can run separate. Lastly, we move to the production cluster.

We have also implemented another limited deployment infrastructure for ISE (see Figure 2). This ISE VM infrastructure mirrors the production cluster in global spread, but has fewer VMs to support the cluster. Scalability for our testing and validation of ISE patches and versions was a key driver for this additional limited deployment infrastructure.

We are pointing several key facilities globally (via their WLCs) to present a new SSID that is not our production wireless SSID. Presenting a new SSID allows us to apply new code on this infrastructure and make it available to specific branch and campus buildings without affecting production access for users. Users involved in testing and validation of a new patch, version, or capability, can attach to the new SSID. If there is a significant issue, they can detach and go back to the production SSID to continue their work.

**Figure 2.** ISE Limited Deployment Cluster



### ISE 1.3 Limited Availability Testing

While Cisco IT may have local access to the business unit creating ISE, our participation in the ISE 1.3 Limited Availability (LA) is similar to a customer applying and being accepted into the program. We have many capabilities queued to launch in fiscal-year 2015, which started in August 2014 and ends July 2015. Our plans for deploying and testing ISE 1.3 LA are extensive.

We are running ISE 1.3 LA as a separate project to ensure that we capture entry and exit criteria. The key features being tested are guest services (new user Interface, common portal, and Application Programming Interface, or API, changes); Endpoint Protection Service (EPS) and re-IP of Network Access Devices (NADs) operational automation, administration and scalability testing, and onboarding (EAP-TLS). Stretch goals for testing are pxGrid and Mobile Device Management (MDM) integration with a third-party MDM solution at Cisco.

---

## Guest Networking and Enhancements

Guest wireless access control was the first ISE capability we deployed globally. Our guest networking system was nearing end of life. The complex system was not very flexible and prone to outages. Simplicity was paramount for our guest networking services via ISE. We required only two servers (primary in San Jose data center, secondary in Western Europe). Having only two servers has lowered our TCO and operational overhead.

Redundancy is built into the application in ISE. This functionality was put to the test in real life when our secondary server failed due to a VM issue. All global guest networking requests landed on the primary server, and all were handled correctly without discernible effect to the end users.

Since December 2014, the average number of guests per month is 60,000 and growing. The tool now has several improvements for guest networking that are being leveraged, including:

- Suspension of guest removes the guest from the network
- Mobile templates (a mobile-optimized template to make login/AUP acceptance on mobile devices much easier)
- Localization of text for onboarding
- Login screens sized for mobile devices
- Notification of the accounts sent to guests via email or SMS

We have made enhancements to the guest services offering with ISE. At Cisco, the policy for guest access is sponsor-based. This means that guests visiting a Cisco facility are required to have an employee create a guest account and provide them with a username and password for access. Within ISE 1.2, Patch 5 was a guest API that allowed other applications to interact with ISE guest capability. We leveraged this feature by linking the badging system used in Cisco buildings, where a receptionist is located on site, to the ISE guest services. When a guest arrives and receives their badge for access to the building, the badge software makes a call, via the API, to the guest services on ISE to populate the required information in ISE for a guest account. The information is sent back to the badging system with the password for access. Guests receive two badges: one for their person to show that they have received proper access to the premises, and another with information required for login. This enhancement reduces the steps for the guest and employee sponsor.

We look forward to more APIs that are scheduled to be included in ISE 1.3, which will further enhance the guest experience.

## Profiling 802.1X Monitor Mode

Profiling is used to classify and manage devices that are attempting to access the network. This ISE capability profiles devices at the network edge using sensing features embedded in Cisco switches, WLCs, and wireless access points. The ISE Profiler classifies devices using the most appropriate endpoint profile, which is configured within ISE and specifies the endpoint identity group. The identity group can be used when defining conditions for authentication and authorization policies.

As a step toward 802.1X enforcement, you need a clear picture of what is on your network. You need to know what devices are capable of doing 802.1X, what devices are not, and where they are located. Then, you can assess the user experience changes you might have to make as a result. Profiling, in conjunction with monitor mode, allows IT to understand what devices are failing authentication and take appropriate action.

---

For profiling, Cisco IT uses RADIUS, Simple Network Management Protocol (SNMP), Dynamic Host Configuration Protocol (DHCP) Helper, Domain Name System (DNS), and HTTP. With the ISE Profiler, we can:

- Perform automatic endpoint profiling
- Provide endpoint behavior monitoring
- Build an endpoint repository without Cisco Network Admission Control

### **Wireless Authentication**

In February 2014, we began deploying ISE to perform authentication for wireless endpoints globally. The initial plans called for three limited deployments (LDs) of ever-increasing size to ensure our methodology was sound and that all issues encountered in the LDs could be addressed before the larger global deployment. We completed the last LD of wireless authentication (from Access Control System, or ACS, to ISE) in May 2014. For the final LD, our target was to get total endpoint count to 50 percent of known endpoints (from ACS figures). The previous LDs focused on a wide range of global, large sites (by endpoints) so that we would get to our 50 percent target. Once achieved, the remaining 50 percent were migrated rapidly.

The general deployment to the remaining 50 percent of the global wireless infrastructure went very smoothly, largely due to the lessons learned in the LDs. The project team has implemented several operational processes to make the deployment and support easier. We have an existing ACS infrastructure that moved the wireless authentication from ACS to ISE. On the WLCs, we made the first three authentication servers the ISE nodes and the remaining two are pointing to ACS nodes. If the ISE system has issues authenticating, it can fail to the first three ISE nodes, and the rest goes to the ACS infrastructure. Additionally, we enabled an analytics report nightly, which searches for WLCs that switch from ISE to ACS. The operations team reviews the report and troubleshoots instances of failover. After there are no failovers for one month, the team will run an automated script to remove the ACS references from the WLCs. All authentication will be done solely by ISE with no fallback to ACS needed. The team that manages the ACS infrastructure can then make decisions regarding the number of ACS VMs needed in future implementations.


Cisco IT will continue to use ACS for numerous tasks, such as TACACS+ for device control, until ISE is able to perform these capabilities. However, as ISE takes over these capabilities, there will be instances when both systems will run in tandem until we are ready to power down the systems that ISE assumes.

### **Replication, Scaling, and Cluster Sizing**

When sizing clusters for devices in our ISE deployment, Cisco IT uses a “3+1” formula: For every person we assume 3 devices (laptop, smartphone, and a tablet) plus one device in the background (security camera, printer, network access device, etc.). In a company the size of Cisco, with roughly 80,000 employees, that equates to 320,000 possible device authentications at once, although the global distribution makes it unlikely that all devices will authenticate simultaneously. With this in mind, the initial sizing was a 4-cluster environment, which we refer to as an “ISE cube” (see Figure 3).



**Figure 3.** Original Multiple ISE Cube



Originally, there were five main clusters, including a U.S. Central cluster. We started the design document in 2011, and focused on the capabilities and scaling in ISE 1.0. The maximum devices ISE could handle in this version were 100,000. Taking future capability rollouts into consideration, we assumed that the clusters needed to handle global deployments of 802.1X for Wired and Wireless; hence, we settled on four production clusters (roughly 320,000 devices divided by 100,000 per cluster). Some regions, such as the U.S., needed multiple clusters because of the large user/device base. We would start with four clusters and adjust as ISE could handle more endpoints. As the product roadmap matured, our plan was to further consolidate the U.S. East and West clusters into one, for example.

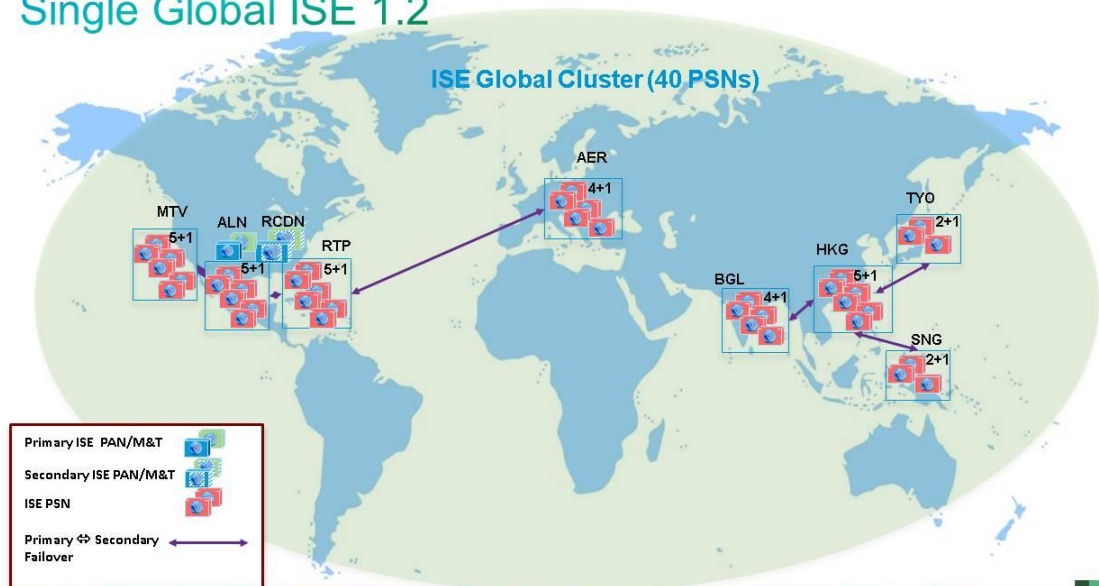
Much of this deployment was a foundation for the IoE. Early discussions around ISE infrastructure sizing focused on the need for us to build the infrastructure so that everything on our network could be authorized and gain access.

We have deployed guest networking, 802.1X Monitor Mode, Profiling globally and are in the process of completing the 802.1X for Wireless deployment. As these capabilities were rolled out into production, they passed from project support to operational support. The team that supports this infrastructure also supports our ACS in production. The experienced team provides valuable insight into production support concerns.

The first official deployment of ISE was on version ISE 1.1.3, with plans to upgrade to the latest patches and then to ISE 1.2. Supporting this global cluster configuration added overhead to the team and also created concerns with regards to patch and upgrade management. The operations team and project team agreed that we could approach sizing differently for the current and short-term capabilities slated for rollout into late 2014. This approach led to a phased uptick in devices on the network that would reach the upper limit. In addition, we had deployed ISE 1.2 with an upper limit of 250,000 devices. Because we had not completely deployed 802.1X Wired globally, only to some high-risk environments where we had intellectual property concerns, we did not need the size to accommodate the full 3+1. Instead, we moved toward a single ISE global cluster (see Figure 4).

Figure 4. Single Global ISE version 1.2

## Single Global ISE 1.2



The change to a single cluster allows us to not have to wait for the Manager of Managers (MoM) expected in later versions of ISE. Multiple clusters mean logging into separate management and administration consoles. Patching and upgrading have to be done at each cluster separately. In current versions, there is no MoM in which these efforts can be centralized. Faced with a static workforce and product limitations that might not exist in the next release, the operational support team decided to merge the four clusters into a single global one.

This one cluster will be sized to the maximum number of policy service nodes (PSNs) allowed in a single deployment (40), with a standby policy administration node (PAN)/M&T pairs to split into two deployments when the limit of 250,000 concurrent endpoints are about to be reached. This split will instantly double the capacity to 500,000 concurrent endpoints without the need to add extra PSNs or change the configuration of the NADs, by leveraging virtual IP addresses (VIPs) and Cisco Application Control Engine load balancers.

At a high level, following are the steps we took to merge the clusters:

- Infrastructure prerequisites:
  - Ensured all clusters were at same version and patch level
  - Disabled all backups
  - Made sure SFTP and ISE servers had enough free storage
  - Completed a backup prior (multiple steps not listed)
  - Ensured primary PAN and M&T VMs were cloned
  - Consolidated all PSN configurations into our development environment
  - Added six new PSN VMs to corresponding VIPs and installed ISE
- Main steps:
  - Merged U.S. West, Asia Pacific, and EMEAR clusters into U.S. East PSN
  - De-registered one PSN behind each VIP in each of the U.S. West, Asia Pacific, and EMEA clusters
  - Registered PSN in U.S. East cluster
  - Added PSN to the other clusters



- Confirmed metrics on U.S. East cluster were correct (now a single global cluster based in our Allen, Texas, data center)
- Checked replication and Active Directory status on cluster hosts
- After migration:
  - Restored services turned off prior
  - Performed a backup
  - Checked all monitoring systems for issues

This work was completed in December 2013, following four weekends upgrading from ISE 1.1.3 to ISE 1.2, one cluster at a time, before all the clusters were merged into one.

We are finalizing the strategy for how to split the deployment in case we reach any of the product limits before the next version. The first option is to split the deployment between east and west, for example U.S. is one deployment, and EMEA and Asia Pacific is another. The second option is to keep the deployments global and split them by leveraging the VIPs to point to a subset of PSNs. This can be a functional split (by service) or a select split (by volume).

Already our support teams have seen the benefits of the single global cluster. Patch 3 and Patch 6 were applied to our global ISE infrastructure of one cluster. This process involved only logging into the PAN and uploading the patch file; the primary manages the allocation of the patch in a serial fashion to each of the secondary servers. Without a single cluster, we would need to apply the patches to each cluster separately. Additionally, with a single cluster, we don't have to worry about versions being out of sync if a patch fails in one of the clusters.

With this consolidation, we have only one-fourth the amount of work to do and a large reduction in risk by upgrading only in one cluster. We have sized the infrastructure for the current demand. Rather than building for all future capabilities as first deployed, we have gone to a right-sizing approach – no excess capacity waiting for future capabilities that are not currently deployed.

We may still get to multiple clusters and, in fact, we probably will. But can't determine at this time when we will need to split the clusters. We do not know the full impact of IoE and all devices that will eventually connect into our enterprise. Using good capacity management techniques is an ongoing effort for any IT organization; the same applies in this instance because the operational team has numerous ways to view and react to changes in load.

A good example of active capacity management is a recent effort to upgrade our VM infrastructure running production ISE. At the time of deployment, we engaged with the internal provisioning team for VMs and selected from a standard set of menu items for this service. After running in production for more than a year, added capabilities were impacting the system. The team noticed large numbers on the disk I/O and memory use. As a result, we upgraded this service in the last three months (see Table 2).

**Table 2.** Current Standards in Production VMs

Hardware	CPUs	Memory	Storage Type	Disk Space (PAN)	Disk Space (MnT)	Disk Space (PSN)	NIC Speed/Count
Recommended Min	4 x 2.4 GHz	16 GB	Dedicated SAN	200 GB	200 GB	100 GB	4 x Integrated GB NICs
Cisco IT PROD	8 x 2.032 GHz (Gold Reservation)	32 GB	Dedicated SAN (except HKG/TYO)	600 GB	1.6 TB	200 GB	1 x 1 GB NIC

Cluster and server sizing decisions are never static. Constant observation and product knowledge have led us to our current optimal deployment configuration. As the capacity is observed, analyzed, and reacted upon, the configuration will change.

### Operational Support

Operational support is split into two main categories: direct support and infrastructure/upper levels of support. Direct support personnel answer calls and perform troubleshooting and escalations. In terms of numbers, the direct support team for the ISE deployment has not changed in the last two years. In Cisco’s fiscal-year 2015, we have not budgeted or forecast any changes in this level of support.

Infrastructure and upper levels of support are handled by the team that manages our ACS and ISE infrastructure. This team was augmented with two additional personnel to handle the new services offered by ISE. In the upcoming fiscal year, we are planning an aggressive capability deployment that includes MDM and expansion of wired authentication globally. We have budgeted to continue with the same two additional Change-the-Business (CtB) personnel. As the capabilities are fully deployed in fiscal-year 2015, this CtB cost will become a Run-the-Business (RtB) expense. Providing this category of infrastructure support is a team that develops and deploys monitors for such services at PEAP, MS-CHAP, and EAP-TLS used by Operations. These personnel are strictly associated with CtB costs and are paid only during their engagement. We have budgeted for one person in this team to be used year over year during the life of the project.

### Challenges and Lessons Learned

Starting in 2012 with ISE 1.1 version, we are continuing our multiyear journey deploying ISE capabilities. Along the way we have compiled several lessons learned.

#### Deploying Guest Wireless Access

- Communication is crucial. This capability involves user experience, so it was critical for us to communicate changes to end-users in multiple ways. Find as many avenues as you can to communicate the changes. This is a case where over-sharing is not a bad trait.
- Design a process for escalations in the first week or two to be handled by a dedicated team. This type of change occurs abruptly for users, so having a special team will ease Day 1 War Room conditions.
- Provide an avenue, whether a forum or mailer, for users to direct how-to questions, provide feedback, or vent their concerns.

### **Deploying 802.1X Monitor Mode and Profiling**

- Consider using a large analytics engine. Cisco IT uses a large commercially available analytics engine to analyze data from ISE. While ISE provides reporting, using a large analytics engine provides additional capabilities.
- Perform platform testing before deployment. Cisco IT has a comprehensive lab with every network device deployed in our global network. In some cases, certain Cisco IOS Software versions did not perform well with ISE features, so we timed the deployment to coincide with version upgrades to ensure the ideal performance.

### **Deploying Wireless and Wired Authentication**

- Identify older AD servers. Older versions of Windows AD servers (2005) on the network can cause intermittent issues with authentication. This is observed when running ACS, not just ISE, for wireless authentication.
- Stage limited deployments. Limited deployments first will help you ensure that platform compatibility is good. What is more, the load on the ISE infrastructure provides the project team the ideal way to manage a global deployment.
- Assess Active Directory server performance in different regions.
- Assess the existing wireless RADIUS authentication log.
- To avoid inconsistent authorization policies, use either PEAP "machine authentication" only or "user authentication" only, not the "machine or user authentication" from Windows supplicant.
- For failover use two Cisco Application Control Engine Virtual IP Addresses (Cisco Application Control Engine VIPs) instead of one VIP in the WLC, even though there are multiple Identity Service Engines behind a VIP already.
- Use logging discriminator to filter unnecessary syslog. Access layer switches are flooded with 802.1X SNMP authentication/authorization syslog.
- Allow UDP port 1645/1656 or 1812/1813 for authentication/accounting traffic from ISE.
- Authentication status is "unknown" after enabling dot1X on access layer switches. This situation requires manual shut/no-shut switch port or manual plug/unplug endpoints.
- Ensure MAB authentication is not permitted for wireless access mechanisms.
- Validate that the versions of IOS running on your L2 devices have no known issues with ISE. There have been instances where memory leaks or other issues are published and finding them before-hand and mitigating them is preferred for any software deployment.

### **Metrics and Communication Steps in Project Planning Process**

- When trying to derive ROI or service metrics, many of the capabilities delivered by ISE will not have a direct cost savings; however, the service metrics are positively impacted as a direct result of the ISE deployment.
- Communicating to end users about changes that will alter the user experience is critical and need to be multichannel. Do not rely on email alone as many users get hundreds of emails daily. Signage, other alert mechanisms, and management support are ideal.

### **ISE Deployment and Support Staffing**

- New ISE support staffing has been required only for the upper levels of support. However, this is not a function of ISE but a function of the capability deployment.
- Leveraging a global engineering staff allows Cisco IT to seek favorable staffing costs, when possible.

- Engage all stakeholders (Architecture, Design, Implementation, Support, Hosting, Storage, etc.) from the beginning of the project.
- Provide hands-on training for implementation and support teams to help ensure smooth deployment and transitions.
- Set expectations about what devices will be 100 percent supported and what devices will be supported but not fully tested.
- Assess the existing ACS-based wireless authentication status.
- Develop an internal Wiki page for the most common end-user supplicants' configuration.

### **Fiscal-Year 2015 Deployment Plans**

Cisco IT has an aggressive capability rollout plan for fiscal-year 2015. We expect to continue deploying Wired 802.1X Monitor Mode and Profiling globally. During fiscal-year 2014, we decided to focus on Wireless Authentication to ISE due to resource and patching on platforms. As mentioned earlier, enabling 802.1X Monitor Mode and Profiling gives us crucial information prior to 802.1X Authentication Mode. As we complete deployments and mitigate issues found on 802.1X Monitor Mode, we will deploy 802.1X Authentication Mode to Wired in those locations.

While we have run EPS (Quarantining) in pilot mode in several key locations, we plan to deploy ISE EPS to all regions deemed as high risk to our intellectual property. Posture Assessment for mobile devices and MDM Device control (via our third-party MDM engine tied to ISE via the API) are also slated for fiscal-year 2015. Lastly, we want to expand our SGA/SGT deployment into labs and for switches where TCAM depletion is at issue.

### **For More Information**

To read additional Cisco IT articles and case studies about a variety of business solutions, visit [Cisco on Cisco: Inside Cisco IT](#).

To view Cisco IT webinars and events about related topics, visit [Cisco on Cisco Webinars & Events](#).

### **Note**

This publication describes how Cisco has benefited from the deployment of its own products. Many factors may have contributed to the results and benefits described. Cisco does not guarantee comparable results elsewhere.

CISCO PROVIDES THIS PUBLICATION AS IS WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING THE IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Some jurisdictions do not allow disclaimer of express or implied warranties, therefore, this disclaimer may not apply to you.




---

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

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

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

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at [www.cisco.com/go/offices](http://www.cisco.com/go/offices).

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