

IT@Intel: Transforming Industrial Manufacturing with Software-Defined Networking

Intel IT delivers improved automation, reliability, scalability, security and resource efficiency to Intel’s factory floor by deploying Cisco Application Centric Infrastructure (ACI) technology

IT@Intel Authors

Chris Benson
Senior Network Engineer

Rob Colby
Principal Engineer

Pat Guerin
Senior Network Engineer

Table of Contents

Executive Summary	1
Limitations of the Classic Factory Floor Network	2
Expanding Our Use of Cisco ACI at Intel.....	4
Benefits of ACI for the Factory Floor Network	5
Results	9
Next Steps.....	9
Conclusion.....	10
Related Content.....	10

Executive Summary

Over the last few years, Intel IT has transformed Intel’s factory data centers using software-defined networking (SDN), leaf/spine topologies and other modernizations. This success has inspired us to apply the same constructs to the factory floor network, which connects and runs the semiconductor tools that produce Intel’s silicon products.

Deploying Cisco Application Centric Infrastructure (ACI) switches powered by the Intel® Xeon® D-1500 and D-1600 processor families has helped enable us to solve some of the technical challenges associated with the classic factory floor network. ACI and SDN bring a wide range of quantifiable benefits, including the following:

- Completed new factory network builds with 85% less headcount and weeks faster through the use of automated scripts.
- Increased efficiency through the use of open-source and vendor-created automation playbooks.
- Implemented a resilient zero-downtime architecture.
- 100x times more efficient on-switch memory utilization for security implementation, compared to the classic network.
- Factory reliability, security, long life and high performance enhanced by Intel Xeon systems on a chip (SoCs).

Traditionally, ACI is used primarily in data centers. We have proven that ACI can meet factory floor requirements and is now used as the plan of record (PoR) for all new Intel factories. We are gradually migrating existing factories to the new network design and anticipate broadening our use of automation through ACI, further enhancing Intel factory efficiency and reliability. We hope that the rigorous testing and validation process that we have established can help other manufacturers transform their own factories using Cisco ACI.

Contributors

- Rajiv Gupta**, Intel Cisco Account Executive, Intel Sales and Marketing
- Gary Morris**, Senior Manufacturing Automation Engineer, Logic Technology Development, Manufacturing Infrastructure
- Ronen Yizhaki**, Manufacturing Network Infrastructure Manager, Intel IT

Acronyms

ACI	Application Centric Infrastructure
ACL	access control list
CLI	command-line interface
DPI	deep packet inspection
EPG	endpoint group
ISSU	in-service software upgrade
PVLAN	private VLAN
SDN	software-defined networking
SVS	Solution Validation Service
TCAM	Ternary Content-Addressable Memory
VLAN	virtual local area network

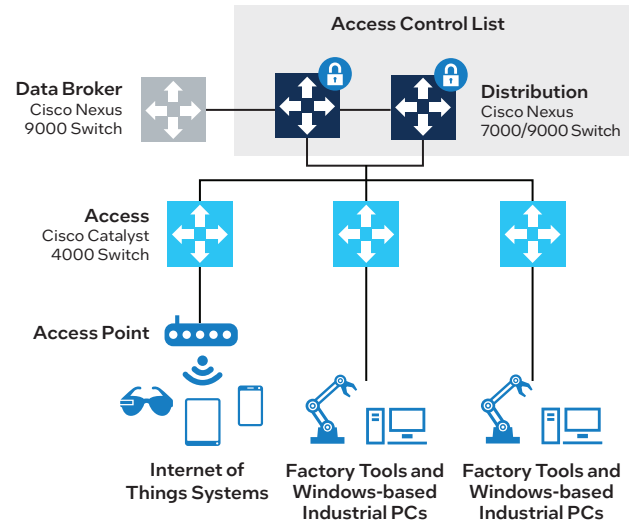


Figure 1. Classic factory floor network design.

Limitations of the Classic Factory Floor Network

One of Intel IT’s top priorities is to help Intel’s factories run as efficiently, reliably and securely as possible. We constantly evaluate and deploy new technologies—like Intel architecture-based industrial PCs, Internet of Things (IoT) sensors and machine-learning applications—to boost quality, yield, security and reliability. A recent area of focus is network transformation for the factory floor network, which provides the factory automation applications, control and data communications for the hundreds of factory process tools involved in producing Intel’s silicon products.

Until recently, we used a classic design for all our factory floor networks (see Figure 1). While the classic design is reliable and functional, it poses several technical challenges, especially because the process tools are provided by a variety of third-party suppliers and require stringent security controls. The challenges of the classic design include the following:

- Limited automation capabilities, which in turn hinder scalability. Manual configuration processes make it difficult to maintain our Copy Exactly! methodology.
- Increased cost and effort for maintaining security due to:
 - Rigid virtual local area network (VLAN) and private VLAN (PVLAN) capabilities, and limited on-switch memory, both of which impede micro-segmentation
 - Inflexible and difficult-to-interpret access control lists (ACLs)
 - Inability to select and redirect specific network traffic for increased inspection
- No support for higher network speeds, such as 100 GbE.
- Some components are reaching end-of-life.

Let’s examine some of these challenges in more detail.

Automation Challenges

The classic network design¹ requires significant manual intervention, which creates a large operational overhead challenge for managing infrastructure. Each switch in the network needs to be configured and managed individually, as a stand-alone component. This configuration and management are accomplished using the device’s command-line interface (CLI), due to a lack of available APIs. In situations where we require automation, we spend substantial time and effort designing robust automation scripts that can interface with the CLI directly. If the CLI syntax changes—as it often does during OS upgrades—these changes can potentially introduce a wide variety of bugs in the automation scripts. Host onboarding is also a manually intensive process, requiring hours of coordination between systems and network operation teams. Combined, these manual processes consume many hours of labor in the factories that use the classic factory network design, and make our Copy Exactly! methodology difficult to synchronize across Intel’s global factory footprint (see sidebar below).

Copy Exactly!

Delivering Manufacturing Solutions Worldwide in Record Time

Copy Exactly! Is a methodology that Intel uses in its manufacturing facilities to transfer production solutions, updates or improvements from one site to another to maintain repeatability, efficiency and reliability for manufacturing semiconductor products. By ensuring that all Intel factories are designed using similar hierarchies and equipment, the Copy Exactly! process minimizes the risk of introducing errors and problems into high-volume manufacturing. Every detail is replicated, including hardware and software components that might affect the manufacturing process.

Security Challenges

Intel’s factories are becoming more complex and increasingly dependent on data, while at the same time our factory production footprint continues to grow—especially as the company expands its foundry business. This growth can complicate how traffic is managed and protected. Three problematic and interrelated areas for the classic factory floor network design include segmentation on-switch memory; ACL readability; and inefficient, difficult-to-scale L4-L7 protection.

Segmentation and On-Switch Memory Limitations

We segment factory process tools to help protect against security threats. Each segment—defined by factors such as supplier, OS type and OS version—has specific policies that limit east/west traffic between tools. The older switches in our factories have limited amounts of Ternary Content-Addressable Memory² (TCAM), which is required to store ACLs. Another issue with TCAM in the classic factory floor network is that the TCAM is inefficiently consolidated in the distribution layer of the network, creating a bottleneck, instead of being spread across all the switches. Increased segmentation creates more complex ACLs, which consume more TCAM. As shown in Figure 2, even a simple ACL can consume over 1% of a switch’s available TCAM. A classic factory floor network may need hundreds of complex ACLs to provide sufficient segmentation and security, requiring us to do one of two things:

- Implement site “TCAM splits,” which involves a labor-intensive effort to manage TCAM usage by spreading TCAM consumption across switches, driving up complexity.
- Purchase new infrastructure to obtain more TCAM, driving up total cost of ownership (equipment and labor).

The limited amount of TCAM also constrains how many PVLANS we can deploy, again preventing us from implementing our preferred granularity of security segmentation. In factories with the classic factory floor network, TCAM was occasionally constrained, and we would conduct weekly, time-consuming TCAM management in an effort to ensure we remain within our planned TCAM utilization range (85% or less of available TCAM).

Cumbersome ACLs

The left side of Figure 2 illustrates sample ACL content.³ This example illustrates that ACLs are not easily interpreted by humans, since they consist of a list of IP addresses and cryptic abbreviations. This type of security framework is not easily readable, manageable or automatable. In addition, it is difficult to obtain a good characterization of the traffic profile flowing through an ACL—that is, which lines in an ACL are being consumed and by what source and destination IPs. If we enable the ACL’s statistics-per-entry capability, which does show this type of information, it consumes even more of the already limited available TCAM.

Classic Access Control List




Figure 2. ACL information from existing switches is hard to read, and even simple ACLs consume a significant amount of on-switch memory.

Inefficient, Difficult-to-Scale L4-L7 Protection

Not all network traffic inspection needs to be treated equally. At any given point, we may want to inspect some traffic deeply and other traffic may be less interesting. However, the classic network design is not flexible or scalable enough to differentiate between traffic that it is important to inspect versus traffic for which inspection is optional. The classic design is either all or nothing—meaning that either all VLAN traffic is connected to our deep packet inspection (DPI) infrastructure and subject to inspection, or the VLAN is not connected to the DPI infrastructure at all, and no traffic is inspected.

Another challenge with the classic design’s Ethernet architecture is that it is impossible to move a specific traffic flow to/from DPI without disrupting the connectivity to the process tool. For example, if a threat arises that causes us to inspect only traffic on port 80 coming from a tool that isn’t already being inspected, we would have to plan a downtime and move all traffic from that network through our DPI solution.

This rigidity means that some traffic that is not required to be routed to the DPI ends up there anyway, driving up costs by consuming capacity and labor. Additionally, as the amount of network traffic that we don’t need to inspect increases, we also run the risk of increased false-positive threat alerts, which can negatively impact the network or process tool and/or require the use of human labor to disposition the false-positive events.

Large Failure Domains

In the classic network design, the access switches to which endpoints (such as factory tools and clients) connect have a large footprint of devices. That is, each access switch could serve a large number of endpoints. As a result, a switch failure could have a large impact, causing unscheduled downtime for a significant number of tools. Scheduled downtime for a switch is also problematic, potentially requesting managed downtime for 100 or more tools. Every change window negatively affects factory output.

Expanding Our Use of Cisco ACI at Intel

The above challenges made it clear that we needed to move to an infrastructure that better supported automation, offered more flexible and dynamic security capabilities, and could reduce the overall impact when planned or unplanned changes occur. The network industry has been trending toward SDN over the last decade, and Intel Manufacturing has been deploying Cisco Application Centric Infrastructure (ACI) and the Cisco Nexus 9500 Series switch on Intel® Xeon® D-1600 systems on a chip (SoCs) in factory on-premises data centers since 2018, gaining experience in the systems and allowing for more market maturity.

In the network industry, ACI is not typically considered for use outside the data center, but we saw an opportunity to take advantage of its capabilities to resolve the technical challenges described earlier in this paper. Therefore, over the last few years we've collaborated with Cisco to establish the suitability of ACI for factory floor process tools, embedded controllers and IoT devices in order to introduce the new technology into the factory environment. Aside from ACI's direct benefits to the factory floor, having a common ACI strategy across manufacturing data centers and the factory floor enables consistency in support and automation systems, since there is a uniform network plan of record (PoR) across the factory.

Introducing Cisco ACI to the Factory Floor

Cisco has substantial experience designing network equipment for the manufacturing environment; however, ACI was developed by Cisco's data center-focused business unit. In 2019, we met with the Cisco ACI business unit's leaders to review and expand on their existing security, scalability and reliability requirements for ACI expansion to the factory floor:

- The network is designed to be highly resilient and focused on enabling the continued movement of Intel's product (silicon wafers) through the factory.
- Intel's factories run 24x7x365, with zero tolerance for unplanned downtime. Planned downtime events for a factory often occur less than one time per year.
- We require high security and micro-segmentation to protect process equipment.

This discussion—over the period of a year—helped position everyone on an ACI for Manufacturing strategy that would address our key challenges and helped ensure that there were no hard barriers. Intel and Cisco jointly collaborated on factory reference designs and various proofs of concept demonstrated that the functionality would meet the needs of the factory floor.

After aligning, we validated that ACI was suitable for both a greenfield startup and that we could upgrade existing brownfield facilities in situ with the least amount of impact.

Helping Build a Highly Secure, Scalable and Flexible Next-Gen Automated Network

Benefits of the Nexus 9300 Series

The Cisco Nexus 9300 Series used in our new Application Centric Infrastructure (ACI) networks is primarily powered by the Intel® Xeon® D-1500 processor family (some SKUs are equipped with the Intel Xeon D-1600 processor family).⁴ Overall, we chose the Cisco Nexus 9300 Series of switches based on Intel architecture because it delivers the following benefits:

Intel® SoC Value

- **Reliability:** Intel SoCs have higher reliability than standard client or server offerings.
- **Security:** Intel® Advanced Encryption Standard New Instructions (Intel® AES-NI) provides integrated support for fast, low-overhead encryption and Intel® Trusted Execution Technology (Intel® TXT) provides platform verification (through authenticated boot) to enable strong security with reduced performance impact. Intel® Secure Key Technology generates high-quality keys for cryptographic protocols.
- **Long life:** Intel SoC lifecycle is between 7 and 12 years.
- **High performance:** We can achieve up to 100 GbE today with the potential of reaching 800 GbE. Other performance-enhancing innovations include Intel® Advanced Vector Extensions 2, Cache Monitoring Technology, Cache Allocation Technology, Memory Bandwidth Monitoring and Intel® Virtualization Technology (Intel® VT).

Cisco Value

- **Flexibility:** These switches support both Cisco ACI and NX-OS. For example, all new factories will use NX-OS for sub-leaves; existing factories will continue to use NX-OS while they gradually migrate to ACI.
- **In-service software upgrade (ISSU):** This feature enables zero downtime during updates.

Greenfield Testing

All functionality was first tested using Cisco’s Solution Validation Service (SVS). Test cases included multipod and multisite, east/west segmentation using group-based-policies, sub-leaf PVLAN to micro endpoint group (μEPG) mapping, in-service software upgrades (ISSU), Intra-EPG isolation/Intra-EPG contracts, Service Graph and policy-based routing to an external security device, CAM⁵ resource optimization validation, L2 extension and DHCP relay. To achieve high-quality testing, we brought online a wide variety of simulated factory clients and tools that were placed on the network, and full negative testing was completed to validate system stability and recovery.

Once full design validation was completed in SVS, these test cases then moved to Intel’s pre-production factory environment to complete an end-to-end pilot project (including proving the benefits of ACI automation). We then moved a large number of our production process tools and associated controllers to ACI, validating that all the tested functionality performed as expected in a live environment.

With our rigorous end-to-end validation process complete, we were then able to execute the industry’s first-ever factory-floor, industrial manufacturing ACI solution. ACI was certified in time for use in new high-volume factories built in Ireland and the U.S. in 2022. Now, ACI will be deployed to all new factories, including factories planned for Arizona, Ohio, New Mexico, Israel, Malaysia, Italy and Germany over the coming years, as part of Intel’s IDM 2.0 strategy (see sidebar). Our experiences and key learnings have paved the way for other manufacturers to gain similar efficiency, security and factory uptime benefits from ACI.

Brownfield Testing

Another critical aspect of our ACI evolution is that we needed to be sure we could upgrade existing facilities to ACI with the least amount of disruption. A full suite of tests was developed, tested and recorded along with the classic PVLAN extension/interaction with new ACI network.

We verified that hosts were working as expected on new μEPGs, as well as validated the proxy Address Resolution Protocol, L3 routing move, and replacing classic ACLs with ACI contracts. Testing was conducted in Cisco’s Labs and with the Cisco Solution Validation service. Once we complete the first brownfield migration (see “Migration Strategy”), all other sites will use the Copy Exactly! methodology to complete the migration.

Benefits of ACI for the Factory Floor Network

Features that informed our choice of Cisco ACI for the new factory floor network (see Figure 3) included the following:

- Product maturity, compared to other SDN offerings.
- High level of automation support, with centralized configuration and the ability to use APIs.
- Flexible micro-segmentation that can scale to Intel’s needs.
- Efficient CAM usage through group-based policy capabilities.⁶
- East/west firewall integration and robust east/west policy controls, including methods of redirection for specific traffic for deeper inspection.
- Built-in fabric-level fault visibility.

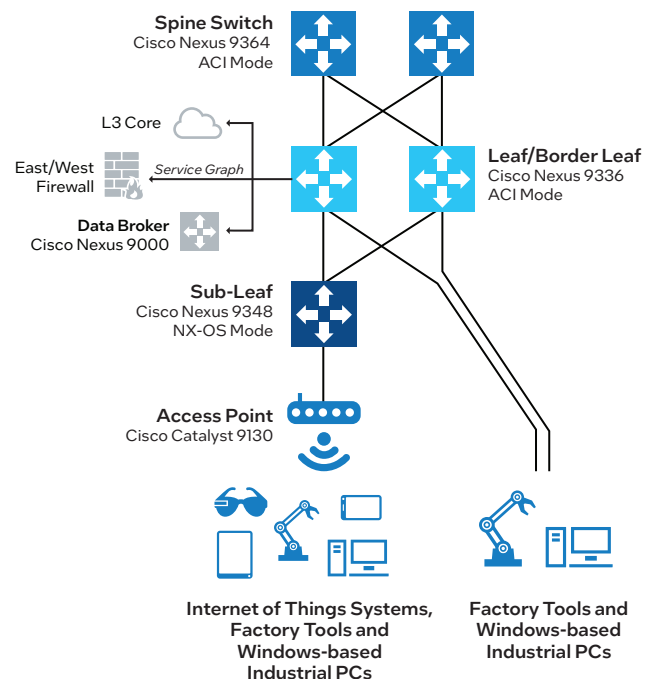


Figure 3. Factory floor ACI network.

The following sections detail how Cisco ACI solves each of the business challenges presented earlier.

Intel’s IDM 2.0 Strategy

Intel Expands Its Plans and Capabilities

Since Intel’s founding in 1968, it has been an integrated device manufacturer (IDM), which means that it designs and builds its own semiconductor chips. In March 2021, Intel’s CEO introduced “IDM 2.0,” a major evolution of that strategy. Intel’s new IDM model includes significant manufacturing expansions, plans for Intel to become a major provider of foundry capacity in the U.S. and Europe to serve customers globally, and expansion of Intel’s use of external foundries for some of its products.

Full Programmability and Automation

ACI is a fully mature SDN solution that supports APIs, scripts and playbooks. The APIs enable us to use an open-source, web-based orchestration GUI. This range of features allows our automation efforts to support all types of users:

- **Non-network-expert consumer.** Our self-service portal framework can provide services such as allowing tool owners to use the GUI to quickly land a tool on the factory floor without involving an IT network expert, accelerating time to production.
- **Network engineer consumers.** Our network engineers are familiar with Cisco products like switches and routers, but have a variety of experience when it comes to writing automation scripts. These engineers can use the GUI to execute our released automation scripts and ensure that our infrastructure releases are implemented per the Copy Exactly! methodology across the environment, without risking human error in applying the configuration. In this way, for example, we can release “automation packages” for landing a new server in a factory data center or for deploying a new type of factory tool.
- **Programmers.** These team members understand both the network and standard manufacturing operating procedures. They write Python scripts and Ansible playbooks that convert typical workflows and configuration processes into automated methods that can be consumed by both the tool owners and network engineers.

As shown in Figure 4, the GUI connects to our centralized production Ansible system, which executes the developed automation scripts. The Python and Ansible back-end automation system scripts are tied to GitHub so that we can ensure quality and integrity through revision control. The GUI abstracts the need for network automation knowledge—just click a few drop-down menu selections to provide a spreadsheet containing specific variables. Using Ansible playbooks and script templates for ACI orchestration helps enforce Copy Exactly! and enables us to achieve repeatability and reduce the operational overhead of bringing a fabric online and making updates.

To enable global consistency, we use a centralized database to hold our global objects, policies, contracts and so on. A repeatable, consistent deployment model that uses the same back-end automation for each install has enabled us to reduce deployment time from weeks to days. To help ensure proper security controls for the automation framework, we have integrated the ACI orchestration workflow with our internal system for enterprise access management.

To accelerate our move to SDN and quickly realize business value from automation, we chose to use open-source Ansible playbooks and vendor-provided scripts from GitHub. Our open market strategy has several benefits:

- There is reduced need to hire expensive dedicated programmers that must understand both network infrastructure and coding practices. Cisco creates an extensive list of playbooks that we can take advantage of, which means we don’t need to develop a high percentage of our automation. However, there will still be situations that require us to modify those playbooks.
- We have access to a code base that is continually being peer reviewed and contributed to by an active automation community.
- We avoid deploying custom code that can be difficult to maintain over time, especially if a programmer leaves the company.

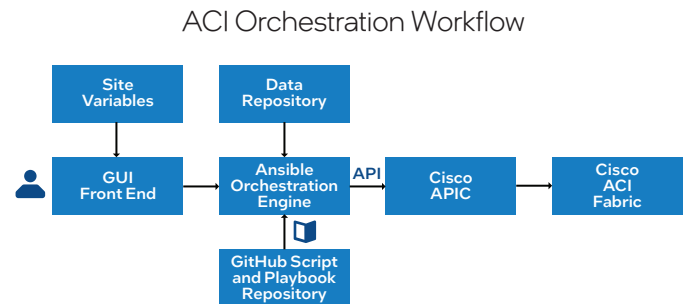


Figure 4. ACI orchestration using Ansible scripts.

To give a tangible example, consider the “new factory automated network build” scripts. The Day 0 process starts with installing the physical network and bringing the Cisco Application Policy Infrastructure Controllers (APICs) online. Then, we are ready to automate the fabric configuration. First, a site-specific template with the local variables (switch names, IP information and so on) is created. Using the orchestration engine GUI, the network engineer links to the Ansible scripts, enters a few required variable values into the GUI and then pushes the variable data to our back-end Ansible/Python script repository to execute the build-out of the fabric across different phases. Examples of phases include fabric discovery; global base configuration; multipod L3 domains; tenant, Virtual Routing Function (VRF) and bridge domains; endpoint groups (EPGs); and security contracts. As the fabric is built out, the APIC is monitored for any faults to ensure a clean installation. The entire process is automated; the engineer only needs to enter some local site-specific data and click Next. The orchestration system executes multiple playbooks in the correct order.

Dynamic and Granular Security Capabilities

ACI resolves many of the security challenges we experience with the classic factory floor network design.

Easier Segmentation and More Efficient CAM Utilization

One of the key factors in choosing ACI for the factory floor was its flexibility in assigning hosts to μ EPGs. The endpoint hosts are grouped to their allocated μ EPGs. ACI gives us the flexibility of automatically assigning hosts to micro-segments, which have a reduced attack surface with a default deny in both the east/west and north/south traffic flows. Contracts (which are similar to ACL controls) are required to allow approved traffic per the security policy associated with each connected system.

Another benefit of ACI is that we can establish an isolation strategy that achieves segmentation regardless of which subnet or VLAN a system resides in. For example, we can segment systems connected to our sub-leaf using a μ EPG that enables us to map a device that’s outside ACI into a set of ACI EPGs and contracts by using device properties such as its IP address and/or MAC address. This enables us to build EPGs and accomplish μ EPG onboarding without having to be specific about the subnet/VLAN association that is required in a classic Ethernet network. All hosts can be on a common large subnet, and through associated attributes we can assign them to different micro-segments. Once process equipment is scanned for viruses and authorized to connect to the network, we use Cisco Port Security to lock down the learned host MAC address to that switch port to prevent any subsequent unauthorized connectivity.

Unlike the TCAM usage problems caused by segmentation with the classic network design, ACI contracts use far less CAM, and the CAM is optimally spread across many leaf switches. Looking back to the example shown in Figure 2, which showed a 1.17% utilization, we can see in the right-hand portion of Figure 5 that the EPG-based contract provides 100x times more CAM efficiency compared to our older switches—only 0.009% with ACI instead of 1.17% with the classic network.

Here is how we estimated CAM utilization for ACI:

$$\text{Number of Contract CAM Entries in ACI} = \text{Filter Entries per Contract} \times \text{Number of Consumer EPGs} \times \text{Number of Provider EPGs} \times 2$$

For an example factory floor, total CAM usage would be as follows:

- $3 \times 50 \times 30 \times 2 = 9,000$ (factory to data center)
- $3 \times 30 \times 3 \times 2 = 540$ (data center to factory)
- Total = 9,540 policy CAM entries/65,536 capacity = 14%⁷

As you can see, there is plenty of room to scale!

Easy-to-Use Security Contracts and Policies

The new architecture enables even better scalability by enabling us to use common contracts to simplify security management. The classic Ethernet configuration includes a complex mix of generic services in every ACL that needs to be written, which is not only an inefficient use of resources but also makes infrastructure scale of common services very difficult to manage. ACI introduces a new method that enables us to apply inherited contracts for common services, increasing efficiency through reusability. What’s more, the group-based policies supported by ACI also remove the IP address element from the contracts, making them more readable and efficient. We are able to arrange systems based on logical groupings that can be easily validated against the expected result, improving audit and validation efficiency as well.

More Flexible L4-L7 Protection

An exciting aspect of our ACI infrastructure is that it brings big improvements for our DPI design. ACI policy-based Redirect Service Graphs support selective traffic redirection when and where it’s needed, without any tool connectivity disruption for redirected flows. For example, we could “on-the-fly” redirect a specific traffic session happening between a factory tool and a server in the data center that are communicating over HTTP to our DPI systems. Even though we have redirected this flow, it will continue as-is without any disruption to the existing TCP connections. This is a significant value considering that we don’t have to ask for tool downtime to make these changes; it’s completely seamless to the endpoint hosts.

Another benefit of this architecture is that we are able to “right-size” the scale of expensive DPI systems to meet the scale of our required inspection traffic. We are no longer captive to an “all-or-nothing” design and can deploy a DPI system that meets the exact scale and size of the traffic we want to inspect at deep levels. For example, we can now have a centralized DPI device that can inspect all critical traffic for a factory site, instead of having to install one DPI instance in each shell.

ACI Group-Based Contract

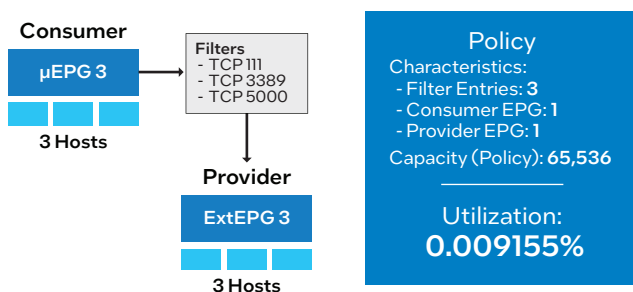


Figure 5. ACI-based security contracts are far more readable than ACLs, and CAM utilization is much lower than the classic network TCAM usage.

Finally, the DPI security appliance can now automatically integrate with our ACI infrastructure using the available APIs for added automation and efficiency. This means we have highly flexible controls that are tied directly to our ACI EPGs and contracts and can control traffic natively in ACI or redirect/copy selected traffic to a security appliance (firewall, intrusion protection system or intrusion detection system) using an ACI service graph. Overall, ACI makes DPI more scalable and enables easier and more secure firewall administration.

Smaller Failure Domains and Less Downtime

The Cisco Nexus 9300 suite of switches offers versatility in product design. Along with implementing SDN, we have transitioned from using chassis-based switches to using multiple low-power, small-form-factor switches to service the factory floor. Careful consideration is given to spreading connected devices across the switches to minimize the impact to the factory operation should a switch interruption or failure occur. The Cisco 45xx switches used in the classic network use 14 rack units, whereas just eight Nexus 9000 switches provide the same port count in eight rack units. Smaller footprint switches allow for easier coordination for refreshes and replacements, with less downtime compared to the types of switches used in the classic factory floor network.

Although there are more switches to manage, we have a common and automated configuration across all switches even down to the VLAN assignment on the host interfaces. Using a combination of a PVLAN and ACI micro-segmentation, we simplify the design by assigning a common VLAN to all ports and using the device IP address or MAC address to assign the host to the correct μ EPG.

Another primary reason for this strategy is that we are able to run our sub-leaf access switches in NX-OS mode, which means we can take advantage of ISSU. A very common yet difficult challenge in the manufacturing industry is to have single-attached tools (as opposed to dual-attached systems in the data center). As mentioned earlier, it is imperative that we minimize impact to a 24x7 operation, where every minute of downtime has revenue impact; this becomes an obvious software upgrade challenge for all of our single-attached devices. ISSU is an important aspect to our strategy of minimizing downtime at the access layer because it allows us to upgrade the system software while the system continues to forward traffic. The upgrades at the access layer can be done with less than one second of interruption to our process tools; we can keep up with security patches and bug fixes while continuing to keep the factory running.

ACI also provides a central view of our infrastructure and lets us more effectively manage the network infrastructure as a whole, instead of managing network devices as individual network components. This avoids the problem posed by the classic network design, where it was difficult to correlate where faults were occurring across multiple stand-alone switches. Overall, ACI has lasting benefits for network visibility, consistency and quality.

Migration Strategy

Intel is now deploying ACI in all greenfield factories that have been coming online, and we will continue to expand use of ACI through Intel's IDM 2.0 build-out. In parallel, we have a strategy to convert our existing brownfield factories to ACI over the coming years. We believe it is beneficial to converge all our factories to ACI quickly to take advantage of the business-enabling and stabilizing features, reduce technical debt and reduce operational burden.

Intel and Cisco have collaborated on the best approach to migrating an existing classic Ethernet network to ACI with the least amount of disruption. The high-level approach to this conversion is as follows (see Figure 6):

Pre-Work

1. Physically build out the ACI fabric in parallel with the existing brownfield network.
2. Establish an L2 connection between the brownfield network and the ACI network.
3. Map out the VLAN-to-EPG association, taking advantage of the ACI feature that allows multiple VLANs to be tagged, and subnets/default gateways to be configured, to the same ACI bridge domain. In this way, we can retain the existing IP addressing with the flexibility to assign VLANs to different μ EPGs. This significantly enhances security because we are not tied to the strict VLAN structure of the classic network and instead can carve up the segmentation per the most optimal security policy.
4. Develop the contracts needed to allow approved communications from/to each EPG.

Migration

1. Negotiate and coordinate downtime with the factory to move the L3 routing from the brownfield (classic) network to the ACI network. When this step is complete, all traffic now goes from a host on the classic network through the ACI network, with all ACI functionality in place.
2. Physically migrate the hosts to the ACI sub-leaf, which can be performed en masse during a change window or per-host outside of a change window on a host-by-host basis.
3. When all hosts are migrated, shut down the classic network.

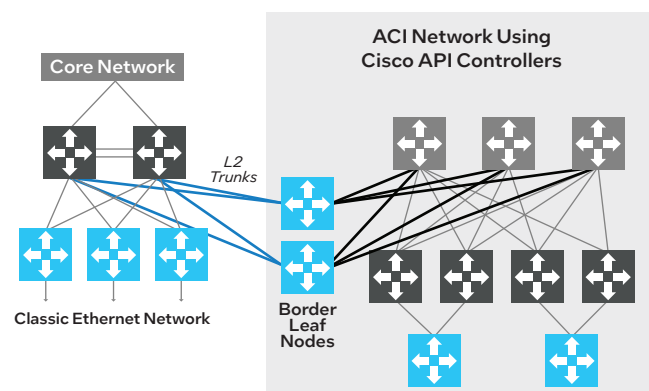


Figure 6. Migration strategy from classic to ACI network.

A Closer Look at the Virtual Factory

Sharing Solutions Across Factories Leads to Increased Manufacturing Efficiency and Quality

Intel implemented a “Virtual Factory” (VF) concept nearly 20 years ago. The foundational assumption is that Intel’s factories have many commonalities. Therefore, sharing solutions and information across all sites helps eliminate unnecessary effort and allows every factory to benefit from a breakthrough solution or idea. Whether it is an ergonomic solution, a new Manufacturing Execution System (MES) or an upgrade to a factory tool, once validated, the change is duplicated using the Copy Exactly! methodology to all the factories. VF also helps ensure product consistency, regardless of the factory from which the product came.

Our AI-based solution that automates gross failure areas (GFAs) on end-of-line wafers is an excellent example of the benefits of VF. We have integrated the solution in the VF network. When the solution finds a new GFA pattern and the yield analysis engineers complete their root cause analysis, the new pattern can be added to the list of known patterns at all factories—improving yield not only at the factory at which the issue was found, but also at all of Intel’s factories around the world.

Results

The following sections quantify and summarize the benefits we are reaping from deploying ACI for our factory floor networks.

Full Programmability and Automation

- New factory network builds completed with 85% less headcount and weeks faster through the use of automated scripts.
- Increased efficiency through the use of open-source and vendor-created automation playbooks.
- Our automation strategy has improved quality of implementation and provides a consistent and standardized environment.
- Our centralized automation framework built on Ansible and GitHub uses open-source packages to accelerate deployment of automation, while taking advantage of a large library of vendor-developed scripts and playbooks to reduce time spent on coding.

Dynamic and Granular Security Capabilities

- We now have a single centralized security policy instead of configuring ACLs and VLAN ACLs on multiple devices. The centralized policy is extended to the fabric edge, not just on distribution switches.
- Security policies are not limited by ACLs to specific networks and hosts on different subnets can be members of the same EPG (that is, intelligent grouping).

- We can create a highly efficient layered security model using multiple predefined inherited contracts for granular security.
- ACI CAM usage is 100x times more efficient than the classic network TCAM usage, and we spend far less time managing CAM utilization for our security implementation.
- More flexible traffic management reduces capacity constraints for expensive DPI equipment and reduces the risk that software bugs on the DPI security appliance will corrupt critical traffic.
- We can create “just-in-time” DPI strategies for any traffic.

Smaller Failure Domains and Less Downtime

- Low-power 1U small form factor switches help minimize the impact to the factory operation should a switch interruption or failure occur and also take up less footprint (8 rack units instead of 14).
- ISSU enables system software upgrades without interrupting traffic forwarding.
- Automation, new switches and ISSU combine to create a resilient, zero-downtime architecture.

Next Steps

Our overarching goal is to automate every aspect of our factory floor network. This won’t happen overnight, but as we make progress, we will generate increasing efficiency benefits. Some specific improvements that we plan to make include the following:

- We will continue to upskill all our network engineers in coding proficiency to maximize the possibilities of ACI API access through scripting, data-driven dashboards and integration with industry-standard systems. This will help improve Day 2 operations in the areas of self-provisioning, monitoring and troubleshooting. The ultimate goal is to increase customer satisfaction and factory uptime as well as reduce deployment time for user requests, such as adding a new tool to the factory floor.
- There are specific needs in the factory space to extend L2 networks across multiple campus shells (either on the same campus or in different geographies). Today, this introduces spanning tree complexity and the need for extra vigilance to avoid loops in the network. In the future, we plan to use the ACI Multi-Site and the Virtual Extensible LAN (VxLAN) protocol to connect multiple fabrics together to extend EPGs and L2 networks between factory campus shells. We are also exploring using ACI Remote Leaf to extend ACI fabrics across different geographies for high-volume cross-site traffic that may not be suitable to traverse firewall boundaries.
- We plan to extend use of ACI to the factories’ Supervisory Control And Data Acquisition (SCADA) system and the Industrial Control System (ICS). Today, SDN is not commonly used with SCADA. But we see opportunities for the business value of converging the network monitoring functions of SCADA and the factory floor network—logical configurations can lead to cost savings. Similarly, applying ACI principles and functionality to the ICS, which operates the factory, can multiply the value of ACI, automation and SDN. Continuing the evolution of the network across all factory functions will multiply the benefits of ACI across currently disparate networks.

- Over time, we plan to take advantage of the Cisco switches' streaming telemetry abilities and implement advanced analytics to increase application visibility. Streaming telemetry will provide better insight into how the factory environment operates. Instead of manually pulling data from the network infrastructure—which increases complexity and doesn't provide full visibility—telemetry will modernize our visibility into the operation of the infrastructure. The new ACI design can seamlessly interact with the intelligence of the system, providing a higher quality of monitoring and alerting. As advanced analytics capabilities mature, we can further take advantage of the data provided by the telemetry, to realize the promise of artificial intelligence and machine learning. We will be able to intelligently interpret the events that happen and even predict things before they happen. This vision will take time to achieve—it's at least two to three years away. But the ability to use network data to proactively detect problems is a powerful motivator.

We understand that network transformation is a journey, and we haven't discovered all the answers and pitfalls. Challenges we anticipate include the following:

- Accelerating ACI global adoption—there is considerable ramp-up time and a learning curve.
- Brownfield migration—even with Copy Exactly! and VF, some individual factories may present unique problems.
- Additional automation hurdles—as Intel's factories evolve with IDM 2.0, we will undoubtedly encounter automation issues that we have not yet anticipated.

Conclusion

Our journey to SDN on the factory floor has introduced many benefits, like the ability to bring a factory network online in a matter of hours instead of days; delivery of consistent builds across factories; and the availability of human-readable security policies that are efficient, highly scalable and automated. With Cisco ACI, we have a scalable method of controlling edge device communication through policies and groups that let us lock systems down to only approved connection flows, increasing our security posture. Better resource optimization is allowing us to bring new capabilities forward, such as micro-segmentation, which has previously been difficult to manage from an operational perspective.

ACI is now an important part of every new Intel factory infrastructure, and we will continue to migrate existing factories to the new network design—advancing automation, enhanced security and smaller failure domains to increase Intel's factory efficiency and reliability. We hope that by sharing our journey of bringing ACI to the factory floor, we can encourage other industrial manufacturers to pursue this same transformation.

Related Content

If you liked this paper, you may also be interested in these related stories:

- [Accelerated Analytics Drives Breakthroughs in Factory Equipment Availability](#) white paper
- [Autonomous Quality in AI Model Productization: A Journey](#) white paper
- [Intel Manufacturing Automation Gets Performance Boost in the VMware vSAN Cache Tier](#) pitch card
- [Developing a Scalable Predictive-Maintenance Architecture](#) white paper
- [Reduce IoT Cost and Enable Scaling through Open Wireless Sensor Networks](#) white paper
- [Driving Improvement in Manufacturing through Advanced Data Analytics](#) white paper

For more information on Intel IT best practices, visit intel.com/IT.

IT@Intel

We connect IT professionals with their IT peers inside Intel. Our IT department solves some of today's most demanding and complex technology issues, and we want to share these lessons directly with our fellow IT professionals in an open peer-to-peer forum.

Our goal is simple: improve efficiency throughout the organization and enhance the business value of IT investments.

Follow us and join the conversation on [Twitter](#) or [LinkedIn](#). Visit us today at intel.com/IT if you would like to learn more.



¹ All new Intel factories use a new, Cisco ACI-based design for the factory floor network; however, the classic factory floor network design is still used in some existing Intel factories. These factories will gradually migrate to the new design over the next few years.

² Ternary Content-Addressable Memory (TCAM) is specialized memory on Cisco switches, often used to store access control lists (ACLs).

³ All IP addresses and port numbers are for illustration only.

⁴ For more information on this family of SoCs, visit <https://www.intel.com/content/dam/www/public/us/en/documents/product-briefs/xeon-d-1500-1600-processors-brief.pdf>.

⁵ In the context of ACI, TCAM is replaced by CAM.

⁶ In ACI, a policy is the combination of an endpoint group (defines what is in a group) and the security contract (defines what communications are allowed for that group).

⁷ Note that the classification into groups based on IP address or MAC is performed in a separate area of TCAM called Longest-Prefix Match (LPM). The default profile capacity for the Cisco Nexus 9300 Series of switches is listed as 64,000 TCAM entries in Cisco's Verified Scalability Guide but the true value is 65,536. The High Dual-Stack Profile has a capacity of 128,000 TCAM entries but does not have as much available LPM as the default profile. Which profile to use varies on switch use, such as client leafs versus border leafs.

Intel technologies may require enabled hardware, software or service activation. No product or component can be absolutely secure. Your costs and results may vary.

© Intel Corporation. All rights reserved. Intel, the Intel logo and other Intel marks are trademarks of Intel Corporation or its subsidiaries.

Other names and brands may be claimed as the property of others. 1222/WWES/KC/PDF