WHITEPAPER

DRAGOS
SAFEGUARDING CIVILIZATION

Network Segmentation Challenges and Solutions

Best Practices for Manufacturers

GLORIA CEDILLO-FARRO | SENIOR INDUSTRIAL CONSULTANT
MIKE HOFFMAN | PRINCIPAL INDUSTRIAL CONSULTANT
DRAGOS, INC.
JUNE 2023

Challenges: What Problems Are We Trying to Solve?

Manufacturers in the Operational Technology (OT) sector face numerous cybersecurity challenges. These challenges arise due to the evolving technology landscape in combination with legacy systems, interconnectivity requirements, and the complexity of equipment and network designs.

Key cybersecurity challenges in the manufacturing industry include:

- **Poor network security controls:** The need for high interconnectivity often leads to the implementation of inadequate network security controls and monitoring practices. This leaves manufacturing plants vulnerable to cyber threats and attacks. 82 percent of manufacturers struggled with poor security perimeters in 2022, according to the Dragos Year in Review.
- **Lack of in-house expertise**: Some manufacturing facilities lack an in-house engineering department and rely on external vendors and integrators for network design and IT support. This reliance can result in challenges in implementing effective cybersecurity measures and network segmentation.
- **Limited visibility and monitoring:** The vast number of connected locations, third-party connections, and remote workers in manufacturing plants make it difficult to achieve comprehensive network visibility and monitoring. This lack of visibility increases the chances of undetected cyber threats. According to the 2022 Dragos Year in Review, a staggering 89 percent of manufacturers had limited or no OT network visibility.
- **OT-specific vulnerabilities:** Operational technology environments share many similar vulnerabilities as their IT counterparts. Yet, OT devices and systems are challenged by long lifecycles, often lack insufficient vulnerability remediation processes, and many of the controllers running the plant floor were not designed with cybersecurity measures in mind. To exacerbate the issue, there is substantial interconnectivity between the new technology and the diversification of assets in manufacturing, including legacy systems.

To address these challenges, implementing cybersecurity controls tailored to the manufacturing industry is critical. This includes adopting recognized cybersecurity frameworks like the SANS Five Critical Controls for ICS Cybersecurity. This framework highlights the importance of developing a defensible architecture, visibility and monitoring, secure remote access, and other important measures.

Threats: What Are We Defending Against?

Manufacturing industries face a wide range of threats, particularly targeting industrial systems that can have indirect but significant impacts on production. If these systems are compromised or taken offline, it can disrupt or completely halt plant operations. Often overlooked components like label printers and their associated servers can also pose risks. Manipulating the label printer server or changing the labels can disrupt manufacturing processes or even stop production altogether.

Ransomware is a major threat to the manufacturing sector as it can affect both the IT environment and production operations. A ransomware attack on a single plant can have cascading effects on the entire supply chain and disrupt interdependent operations in other plants. The manufacturing sector is a prime target for ransomware attacks, with 73 percent of such attacks in recent times directed at manufacturing, according to the Dragos 2022 Year in Review.

Well-known ransomware attacks like WannaCry and LockerGoga have impacted the manufacturing sector's IT environment, indirectly affecting operations in the operational technology (OT) side. These attacks have necessitated switching certain functionalities to manual mode due to extensive interconnectivity between IT and OT systems. Specifically designed ransomware like **EKANS targets industrial control systems (ICS) and OT assets directly.** While it may affect IT systems such as SQL databases, its primary goal is to disrupt or impair critical OT systems and visibility into them.

The OT environment shares similarities with the IT environment, which means adversaries can exploit known vulnerabilities and weaknesses in the network. Unmonitored remote connections can be leveraged to gain unauthorized access to the OT network and potentially compromise control systems. If flat networks exist on the lower levels of the Purdue level, which is common in manufacturing, this means that nodes from different sections can communicate freely without any segmentation. This creates additional opportunities for attacks. Developing a defensible architecture is crucial for the manufacturing sector and is one of the critical recommendations outlined in the five ICS Cybersecurity Critical Controls.

The Five Critical Controls for Manufacturing Cybersecurity – and the Role of Network Segmentation

The SANS Five Critical Controls for ICS Cybersecurity can be applied to manufacturing environments:

1. ICS Incident Response Plan (IRP)

OT incident response is unique and distinct from IT. OT involves different device types, communication protocols, adversarial tactics, techniques, and procedures (TTPs). Containment is done at the device, vlan, or between the IT/OT boundary. Investigation requires a different set of forensic tools and technology constraints. Managing the potential impact of an incident is also different for manufacturing plants. Therefore, OT should have a dedicated IRP created from known scenarios that are impacting the industry, such as ransomware. The IRP should include a definition of an event and criteria for incident declaration, contacts (including vendors), and logical steps from identification through recovery response phases. The IRP should be validated by conducting Tabletop exercises to test and improve response plans.

2. A Defensible Architecture

OT security strategies often start with technical controls at the network, where communications between IT and OT are restricted. However, many manufacturing plants struggle with this control alone. In Dragos 2022 Year in Review study, we found that 82% of manufacturers had poor security perimeters. Beyond IT/OT separation, the OT architecture needs to be designed with least privilege in mind. If an OT device doesn't need to communicate with another device, it should be placed in a restricted zone. By creating zones of trust with conduits of communication, the plant floor can be made much more secure and equipped to handle additional network monitoring controls. Perhaps even more important, however, are the people and processes to maintain these security controls, as they can quickly atrophy over time if not maintained by trained personnel.

3. ICS Network Visibility and Monitoring

Visibility and Monitoring: You can't protect what you can't see – and manufacturers are behind the curve, with 89% having limited or no OT network visibility. A successful OT security posture maintains an inventory of assets, maps vulnerabilities against those assets (and mitigation plans), and actively monitors traffic for potential threats. The visibility gained from monitoring your industrial assets validates the security controls implemented in a defensible architecture. Threat detection from monitoring allows for scaling and automation for large and complex networks. Additionally, monitoring can also identify vulnerabilities easily for further mitigation.

4. Secure Remote Access

Secure Remote Access: Secure remote access is critical to OT environments. One method, multi-factor authentication (MFA), is a rare case of a classic IT control that can be appropriately applied to OT. Implement MFA on Remote Access solutions used for both employees and vendors alike. Where MFA is not possible, consider alternate controls such as jump hosts with focused monitoring. The focus should be placed on connections in and out of the OT network (82% of manufacturers have undocumented or unmonitored external connections into their OT environments) rather than connections inside the network.

5. Risk-Based Vulnerability Management

Risk-Based Vulnerability Management: Knowing your vulnerabilities – and having a plan to manage them – is a critical component to a defensible architecture. While patching an IT system like a worker's laptop is relatively easy, the effects of a patch gone wrong in manufacturing could shut a plant down, resulting in significant production loss. An effective OT vulnerability management program requires timely awareness of key vulnerabilities that apply to the environment, with correct information and risk ratings, as well as alternative mitigation strategies to minimize exposure while continuing to operate.

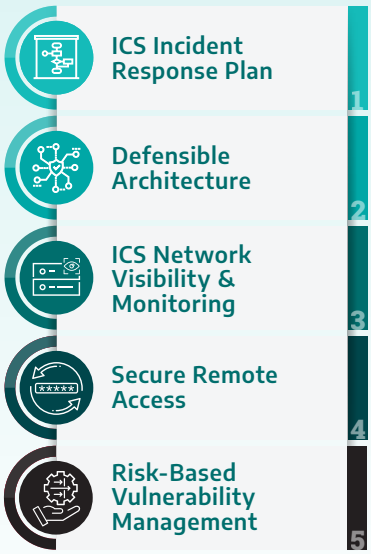

These five critical controls provide a framework for manufacturers to strengthen their cybersecurity posture and mitigate risks in their ICS/OT environments. By implementing these controls, manufacturers can enhance their ability to detect, respond to, and recover from cyber threats and incidents, safeguarding their operations, assets, and data.

In this white paper, we get into detail about network segmentation, which is an important step in achieving a defensible architecture and enhancing visibility and monitoring capabilities. We also explore how segmentation works at the lower levels of the Purdue model, including L2 and L1. Manufacturers should implement network segmentation for several reasons:

1. **Enhanced Security:** Network segmentation helps improve overall security by dividing the network into smaller, isolated segments. Each segment can have its own security controls and access permissions, reducing the attack surface and limiting the potential spread of cyber threats. If a breach or compromise occurs in one segment, it is contained within that segment and doesn't affect the entire network.

2. **Protection of Critical Assets:** By segmenting the network, manufacturers can isolate and protect their critical assets, such as production systems, intellectual property, sensitive data, and control systems. This ensures that even if other parts of the network are compromised, the critical assets remain secure and inaccessible to unauthorized individuals.
3. **Operational Continuity:** Network segmentation can help prevent disruptions in manufacturing operations. If a cyber incident or malware outbreak occurs in one segment, it can be isolated and contained without affecting the production processes in other segments. This helps maintain business continuity and minimizes the impact of potential cyberattacks or system failures.
4. **Granular Access Control:** Network segmentation allows manufacturers to implement granular access controls, ensuring that only authorized individuals or devices can access specific network segments. This reduces the risk of unauthorized access, data breaches, and insider threats.
5. **Improved Network Performance:** Segmentation can also lead to improved network performance including the lower levels of the Purdue model L1 and L2, such as HMIs Clients, servers and Controllers. By segregating traffic and resources into different segments, manufacturers can prioritize critical applications and optimize bandwidth allocation, leading to better network performance and reduced latency.

Network segmentation is a proactive security measure that helps manufacturers protect critical assets, maintain regulatory compliance, ensure operational continuity, and enhance overall network security. It enables organizations to isolate and control access to different parts of the network, reducing the potential impact of cyber threats and providing a more secure and resilient infrastructure.

Segmentation Using Visibility

In manufacturing, maintaining a segmented architecture is a significant challenge due to the diverse range of technologies and connectivity requirements. The absence of segmented architecture complicates the implementation of traffic monitoring across multiple plants. In many cases, the lower levels of the Purdue Model are left unmonitored because of flat architectures that lack sufficient managed switches capable of supporting mirrored ports. Flat networks lead to a lack of designated chokepoints. Without these choke points, plants may remain unaware of substantial amounts of network traffic from the plant floor. As a result, they may need to incorporate additional sensors and monitoring equipment, leading to increased visibility costs. To provide further context, the Dragos Platform Asset Map is shown to illustrate how adding visibility to the network can also drive segmentation improvements, including at the lower levels of the Purdue model.

Visualizing Improvements in Segmentation – The Demilitarized Zone (DMZ)

The illustrated scenario in Figure 1 begins with establishing fundamental segmentation between corporate and OT networks, which is often an initial step in implementing segmentation within manufacturing environments.

Figure 1 reveals a notable absence of a demilitarized zone (DMZ) between the traffic originating from the corporate network and entering the OT environment. This lack of a DMZ escalates the risk of permitting untrusted traffic into the OT environment. A DMZ can act as a controlled connection point between corporate and OT systems by terminating connections from the IT side and establishing new connections to access the OT network.

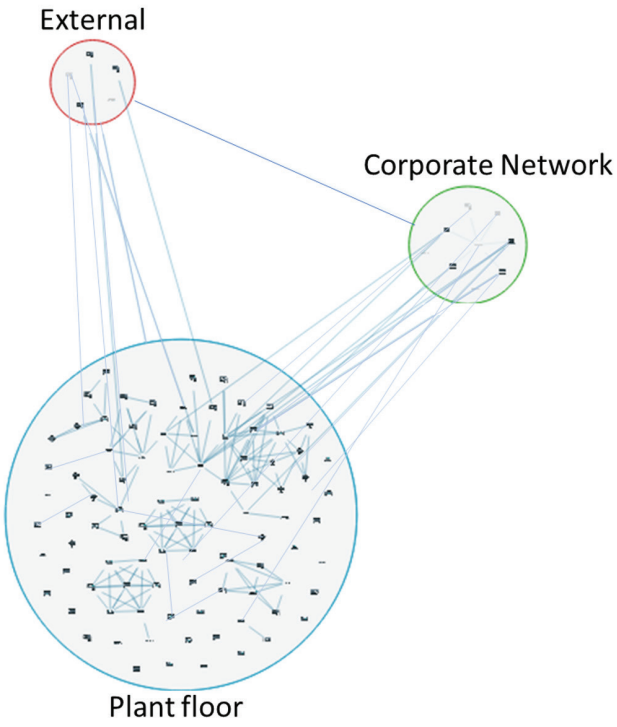

Figure 1: Original Architecture

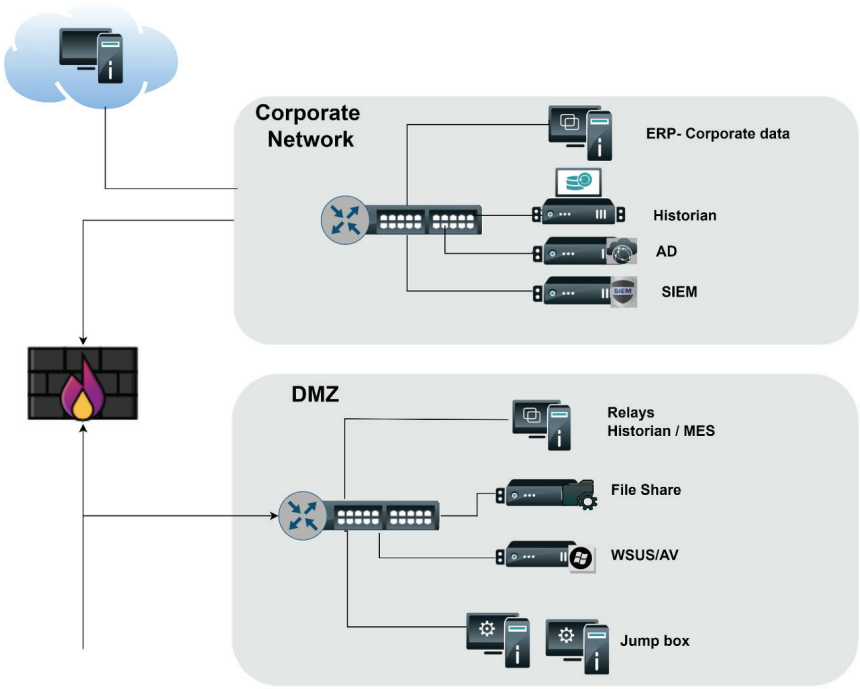

Figure 2: Network Diagram with Addition of a DMZ

The DMZ also handles remote connections and monitoring tasks. Within the DMZ, various common activities take place, including jumping hosts, remote connectivity, file transfers, patching, monitoring solutions, and other relevant tasks. Each of these activities is further divided into distinct "micro" VLANs per service for segmentation purposes. **Figure 2** illustrates a network diagram with the addition of a DMZ.

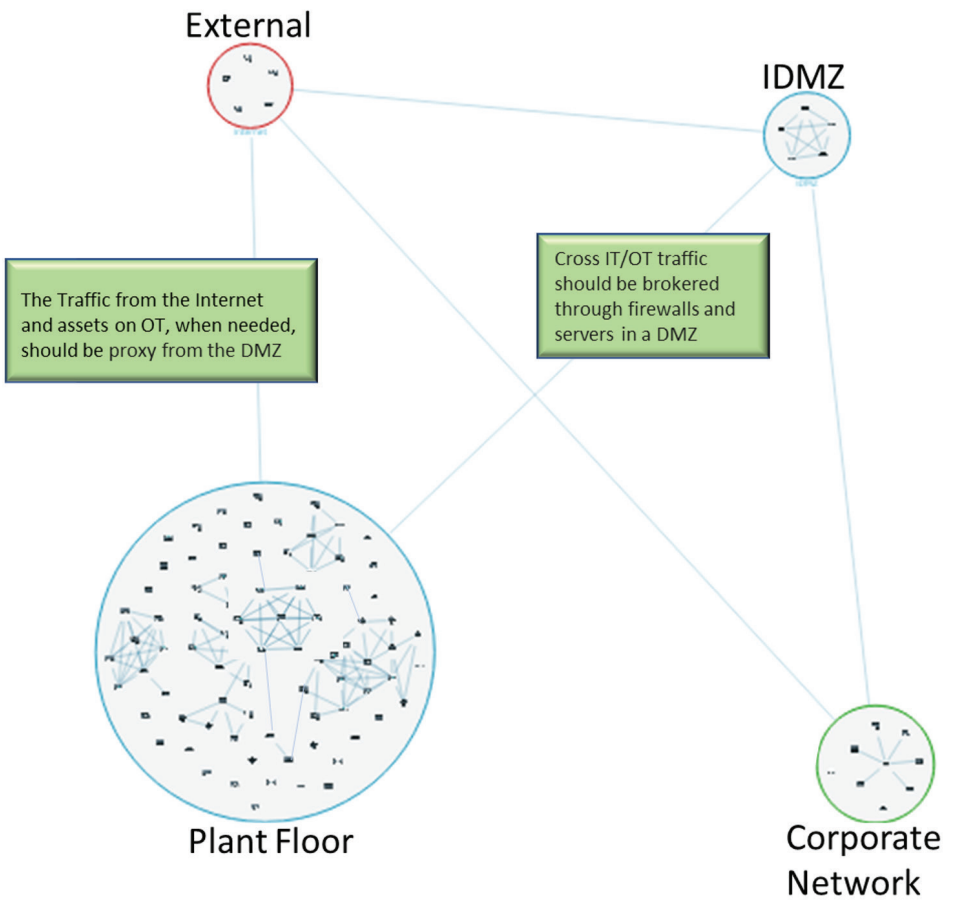

Figure 3: Dragos Platform Assets Map, Incorporating a DMZ

By implementing a demilitarized zone (DMZ), the volume of traffic can be reduced, limiting the adversary's ability to establish a connection from corporate assets to the OT network. However, in **Figure 3**, it is evident that the OT assets remain interconnected within the same network and across the internal plant. This setup presents challenges for monitoring the OT network and adversely affects its performance in automation. Consequently, it becomes necessary to introduce segmentation at different levels between control assets, supervisors, and management.

Visualizing Improvements in Segmentation – Level 3 Enhancements

The L3 section should consist of assets responsible for data interchange across L2 systems and between the corporate and OT networks. The key distinction between DMZ assets and the L3 section lies in their respective roles. DMZ assets primarily focus on providing security and access control to the OT network, while L3 assets facilitate the transfer of authorized data from IT to OT through the DMZ. **Figure 4** illustrates this concept, where the DMZ encompasses the Manufacturing Execution Systems (MES) relay.

This relay acts as the intermediary point for transferring production data to Enterprise Resource Planning (ERP) and Business Planning and Logistics systems. This architecture ensures that ERP systems do not have direct access to L3 MES servers for information and data exchange, thereby enhancing security and control.

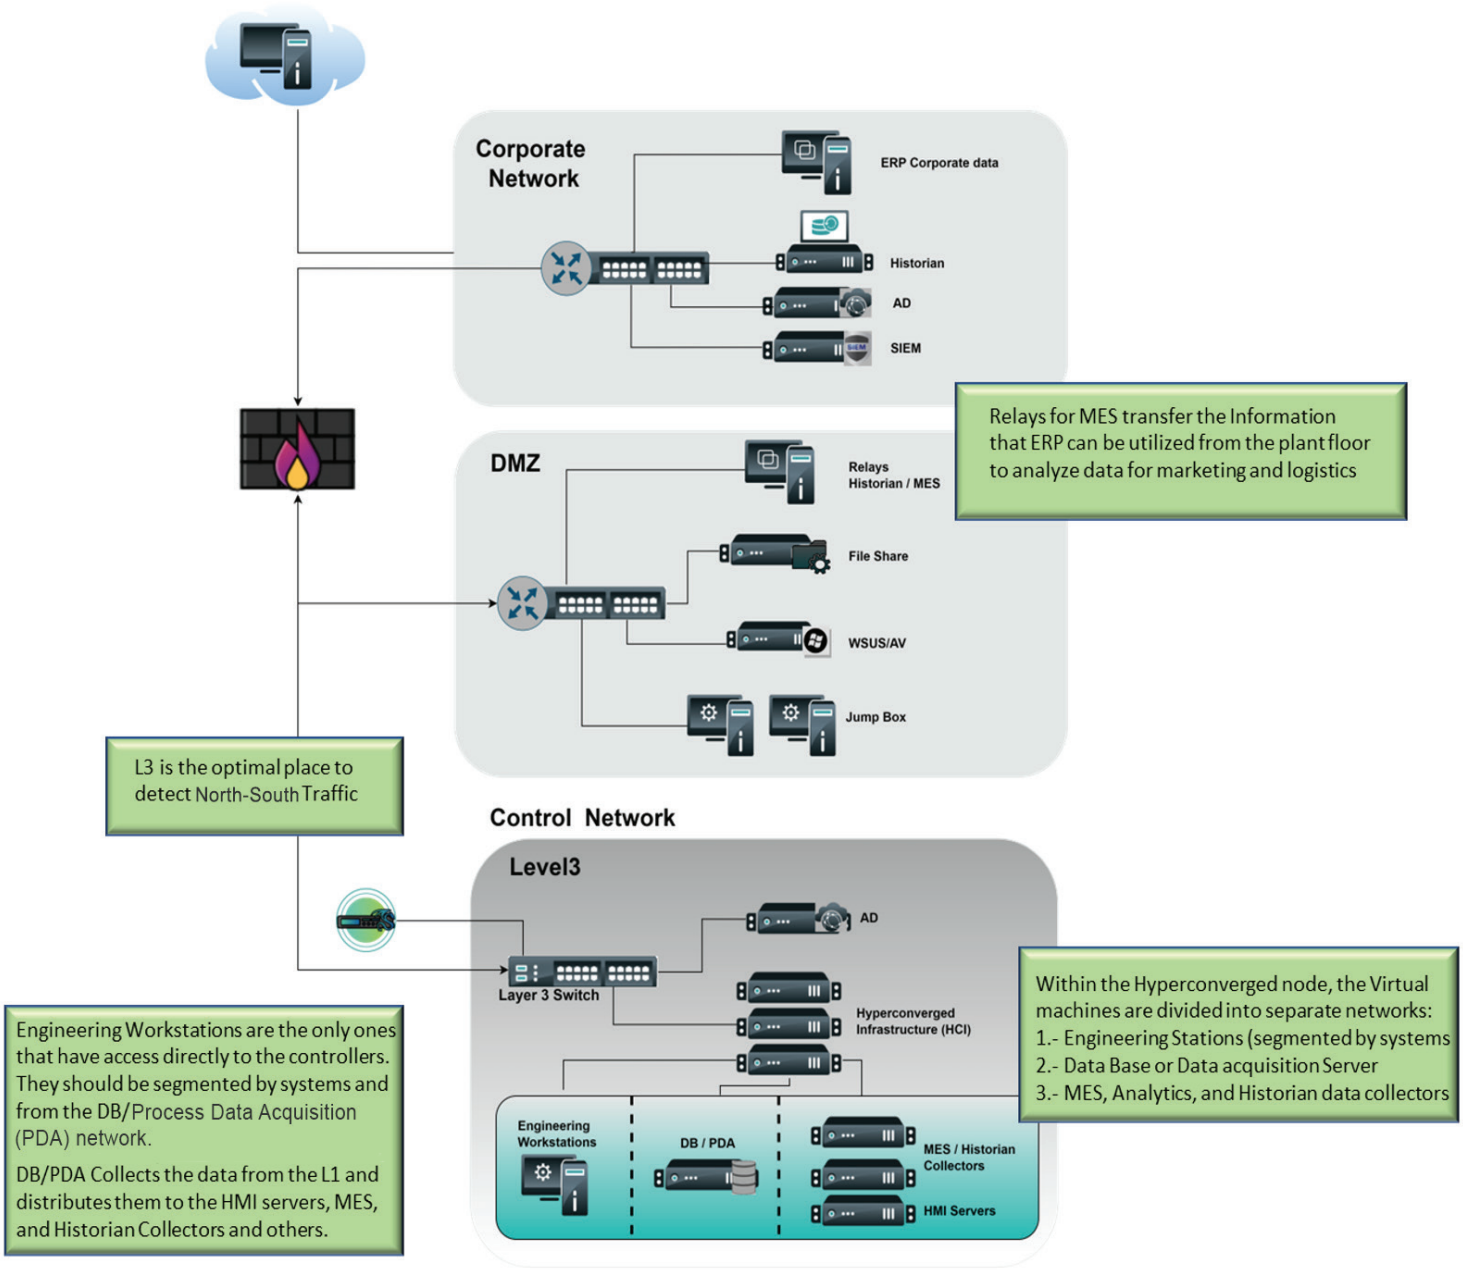

Figure 4: Network Diagram Adding Level 3 Segmentation

ERP AND MES DATA FLOWS

Here are a few examples of the data exchanged between Enterprise Resource Planning (ERP) systems and Manufacturing Execution Systems (MES) during operations:

- **ERP to MES:**
  - Batch sizes, assignments, and start times
  - Setpoints, recipes, and constraints
- **MES to ERP:**
  - End times, outages, quality parameters, and performance times
  - Progress and equipment availability

From this list, it becomes evident that ERP systems are responsible for managing business and logistic systems, while MES systems handle data interface within the operational systems.

Note: The ISA95 standard provides guidance on the implementation of data flow between these two systems.

The L3 segment also accommodates assets used for network configuration, including application servers, Active Directory, and administrative function applications. It has been observed in various industries that a single Active Directory domain is often employed to encompass all users and assets across both IT and OT networks. However, this configuration poses a significant security risk.

Having a single Active Directory domain across multiple security zones creates an opportunity for adversaries who gain user access to traverse through multiple security zones. If there is only one domain that encompasses IT, the DMZ, and OT systems, the barrier for unauthorized access to propagate across the network becomes considerably lower compared to a scenario where an adversary would need to compromise a separate Active Directory environment shielded by additional firewalls and security controls.

To enhance security in an environment where multiple systems require data from controllers and plant floor systems, it is essential to establish a network architecture that incorporates a Process Data Acquisition (PDA) or Database (DB) server. These servers act as intermediaries between the controllers and other data-consuming systems such as HMI servers, MES, analytics, and historians. This arrangement enables the implementation of robust security controls, including more precise Access Control List (ACL) rules within a firewall. Additionally, it facilitates segmentation and direct access restrictions for assets at the lower levels of the Purdue Model.

Engineering Workstations (EWS) are also considered part of L3 (in most cases) due to their functionality. However, since they directly communicate with assets at lower levels of the Purdue Model for configuration and controller logic changes, it is advisable to place them in a separate VLAN or network compared to the PDA/DB server. This segregation aims to impede lateral movement for adversaries who gain direct access to the control network at

lower levels. As mentioned earlier, these assets have direct connectivity with most of the assets at the lower level of the Purdue Model. The Level 1 section further emphasizes the importance of this separation as data interchange communication is facilitated through a designated "MAIN" PLCs.

Virtualization is predominantly observed at this level as it offers cost savings, scalability, and, in some cases, high availability depending on the chosen hypervisor architecture. **Figure 4** provides an example of a hyperconverged node that handles storage, network, and virtualization. The selection of a virtualization system depends on factors such as plant size, containerized systems, and scalability requirements. However, it's crucial to note that when choosing a virtualization system, the availability requirements should align with those of the underlying control system.

In terms of visibility, the core switch at Level 3 serves as a significant choke point. It acts as the convergence point between OT zones, and the traffic flowing up to and out of the DMZ often passes through this section of the switching fabric. Consequently, capturing and monitoring the traffic at this point allows for the detection of any unauthorized or abnormal communications occurring across multiple zones.

Figure 5 showcases how the Dragos Platform identifies the new zones, providing a visual representation of the detection process.

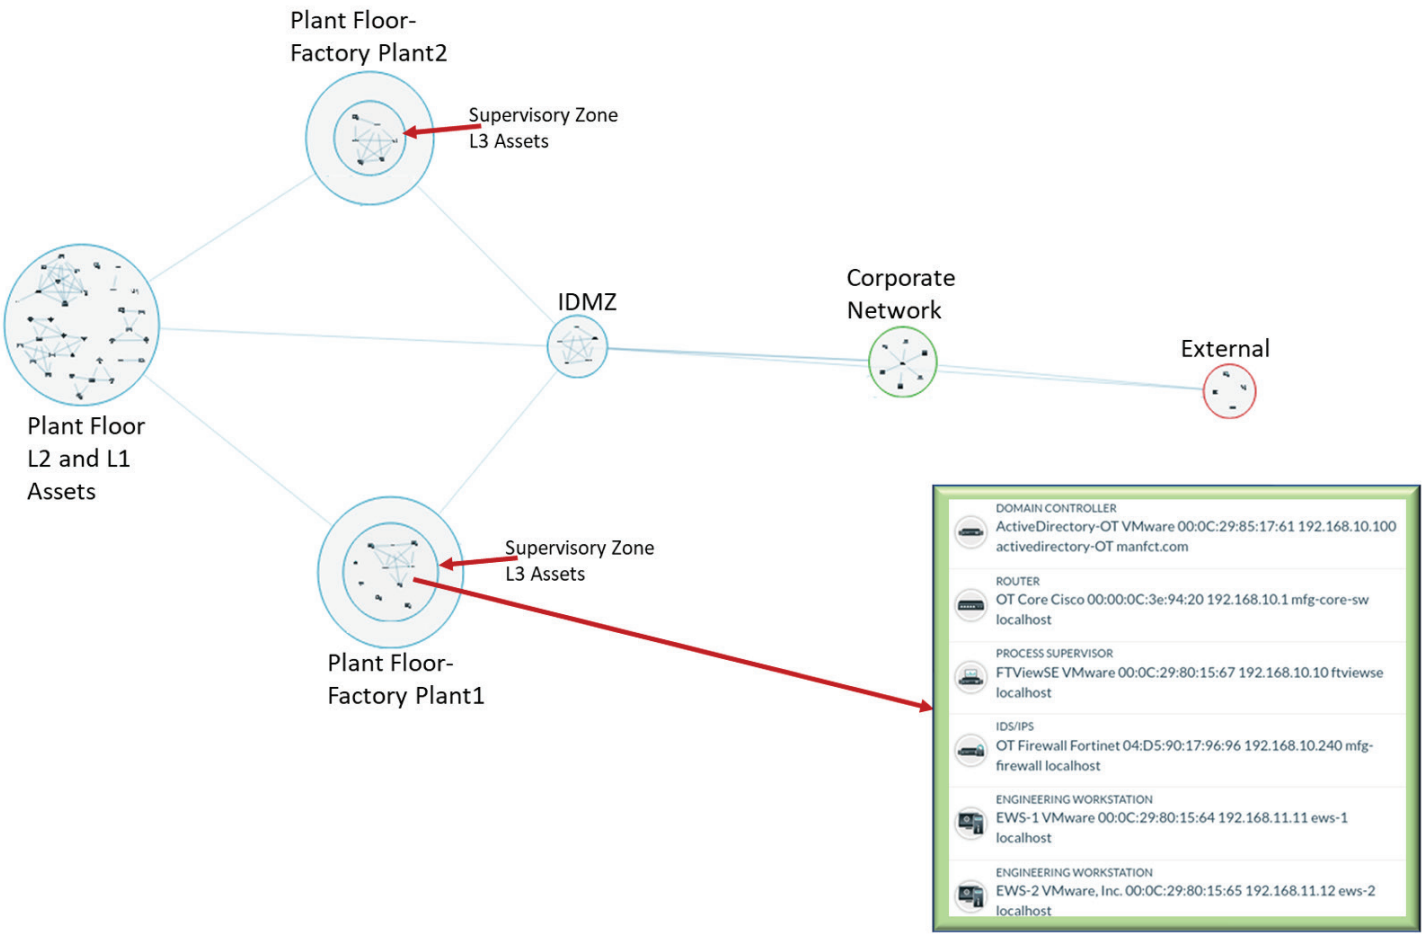

Figure 5: Dragos Platform Asset Map with L3 Zone

Visualizing Improvements in Segmentation – Level 2 Enhancements

Level 2 assets primarily consist of HMIs (Human-Machine Interfaces) and servers used for operating and monitoring plant processes. HMIs can be categorized as Operation Panels (OP), Electronic Operator Interfaces (EOI), and HMI clients. OP and EOI are often used interchangeably for local operations. As depicted in **Figure 6**, these local HMIs establish direct communication with the plant controllers. However, it is highly recommended to separate them from the main controller network. Additionally, operation panels should only communicate with the specific controller(s) they locally operate and monitor. For instance, if an operation panel oversees three conveyors controlled by two PLCs, it should exclusively communicate with these two PLCs through a separate VLAN distinct from the main controller network. Conversely, the main controller network serves as the communication hub for all PLCs within the area, enabling inter-PLC communication.

On the other hand, HMI clients are assets that centralize the operation of different system sections or the entire operation. These clients are typically situated on control pulpits or control rooms, depending on the system's design. **Figure 7** illustrates that the HMI client interface runs through the HMI servers, and thus, they do not directly communicate with the PLCs. Therefore, HMI clients should always be connected to a separate network.

A common mistake seen in the industry with HMI servers is attempting to centralize HMI clients into a single HMI server, even when multiple distinct or different integrated systems exist within a line. Despite being

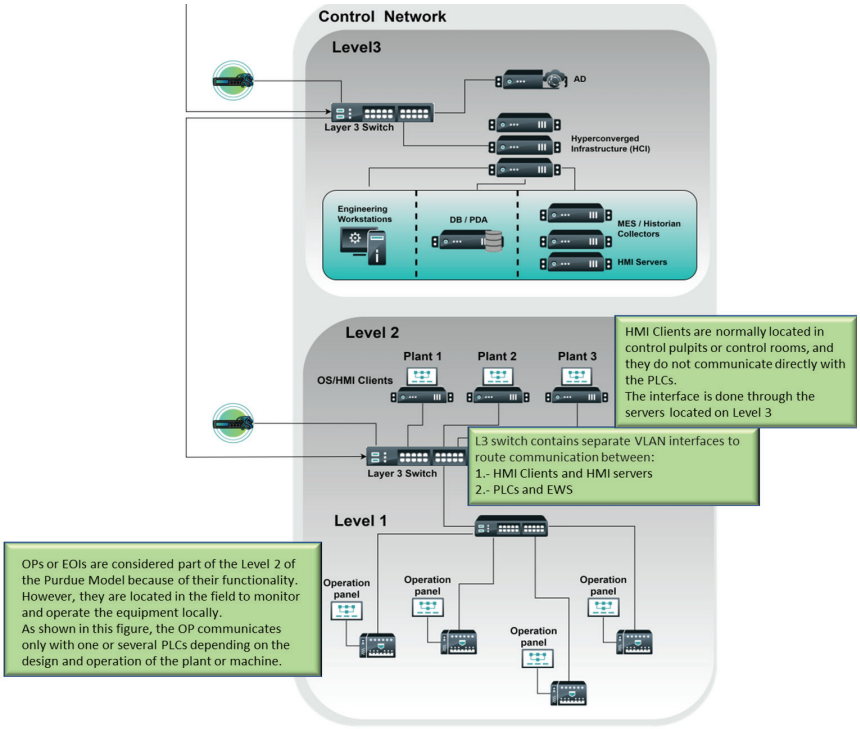

Figure 6: Network Diagram Adding Level 2 Segmentation

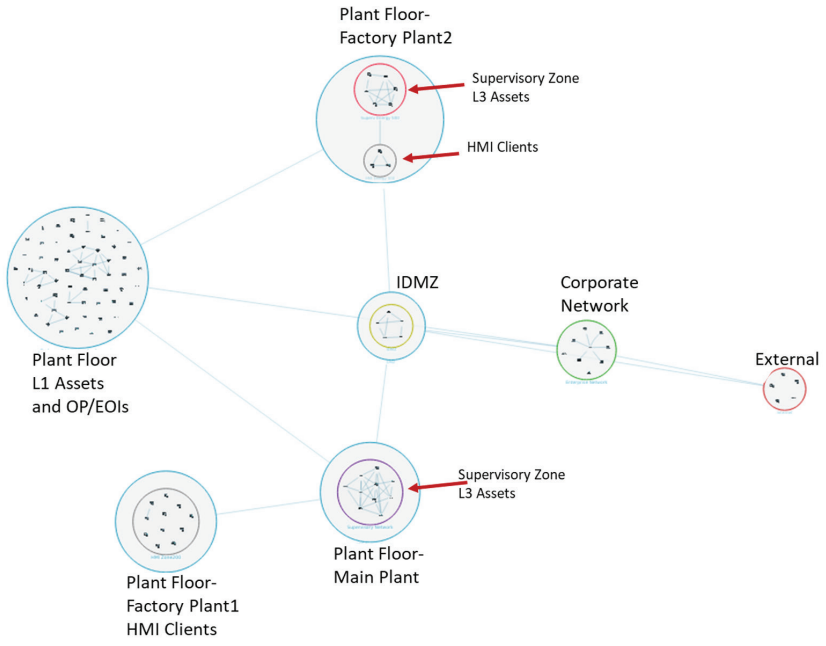

Figure 7: Dragos Platform Assets Map with HMI Clients

in the same unit, their operations and monitoring are independent. In such cases, it is more effective to separate servers by operational sections, allowing for distributed functionality and better control of HMI clients and controller interfaces. It is even recommended to have separate HMI servers when different plants or lines are present. For example, if multiple assembly lines producing different products do not require operation or monitoring by the same HMI clients, each plant should have its own clients with their respective HMI servers. This type of architecture begins with dividing zones from high to low based on trust levels. Segmentation then facilitates the implementation of security controls that restrict interfaces between assets where communication is necessary.

The OSI Layer 3 switch connecting Level 2 and Level 1 systems, shown in **Figure 6**, is another choke point to consider using for monitoring as it often contains traffic between Level 3, Level 2, and Level 1 systems. Monitoring this east/west traffic provides significant visibility into underlying controller data, commands and controller changes coming from HMIs, servers, and engineering workstations.

Visualizing Improvements in Segmentation – Level 1 Enhancements

Level 1 in the Purdue Model is dedicated to hosting system controllers, including PLCs, remote I/O, variable frequency drives (VFDs), smart motor starters, and related assets. Typically, the architecture at this level is flat, as depicted in **Figure 7**, where all controllers are interconnected within the same network. This allows them to communicate with one another, although data transfer requires intentional configuration of telegrams and communication channels. However, it's important to note that communication and access to each asset are possible. In scenarios where different systems, plants, or lines coexist, it is highly recommended to implement network separation through zones, and unnecessary communication between controllers should be avoided.

Like the PDA server structure, if multiple systems with several controllers need to communicate with other systems, it is advisable to employ a centralized interface using a communication PLC as a broker. This communication PLC is referred to as the "MAIN" PLC in this context. **Figure 8** demonstrates that the MAIN PLC can also serve as the primary interface for communication with other plants if necessary.

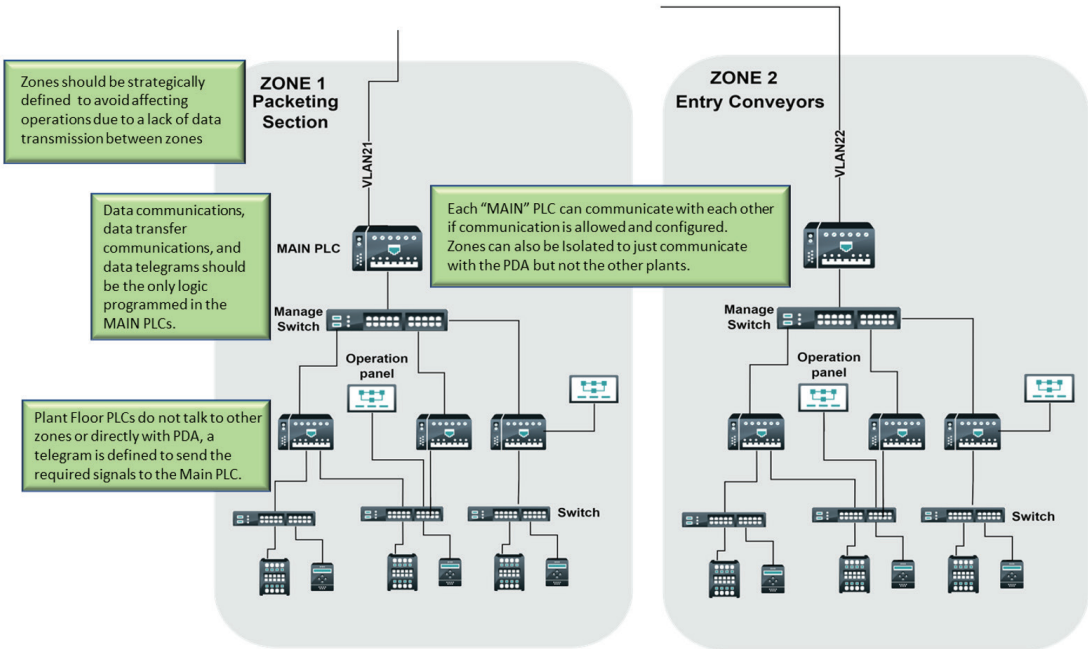

Figure 8: Level 1 Zones Segmentation

This configuration reduces the implementation of firewall ACLs, as the rules are limited to the required communication with the MAIN PLCs. Moreover, this architecture prevents direct communication between the PLCs and HMIs responsible for controlling and monitoring machines at the factory floor level, ensuring their placement in different VLANs.

This architecture brings several benefits, including improved network performance and reduced latency. HMI servers and communication with other plants are channeled through the MAIN assets, which handle data distribution and read/write processes to lower levels. Additionally, this type of architecture, featuring zones and a centralized interface controller, simplifies the process of adding new assets to the network. The addition of a new asset under the MAIN asset only requires an interface and a channel to the respective PLC, along with an optional connection to an ACL rule if a security control is in place between the MAIN and the new PLC. Furthermore, modifications to the logic of a PLC at the factory floor level do not impact the logic of the MAIN PLC, ensuring greater flexibility and ease of maintenance.

The primary disadvantage of having a PLC data concentrator is that if the plant does not have local operations, then losing the MAIN PLC can provoke a loss of view with the HMI clients that hampers the operation of an entire zone. In that case, the plant should consider the risks and evaluate all security controls implemented to decide if redundancy on the MAIN PLC is required.

From a visibility perspective, as shown in **Figure 9**, passively monitoring network traffic can be done at the MAIN PLC asset switch where a sensor can inspect all device traffic and communications between other plants or systems. However, if several MAIN assets exist within a zone, monitoring each of the zones can incur higher costs. Therefore, it is recommended that the plant review critical systems by performing a **Crown Jewel Analysis** to determine which zones are of the highest risk and thus require more detailed levels of visibility.

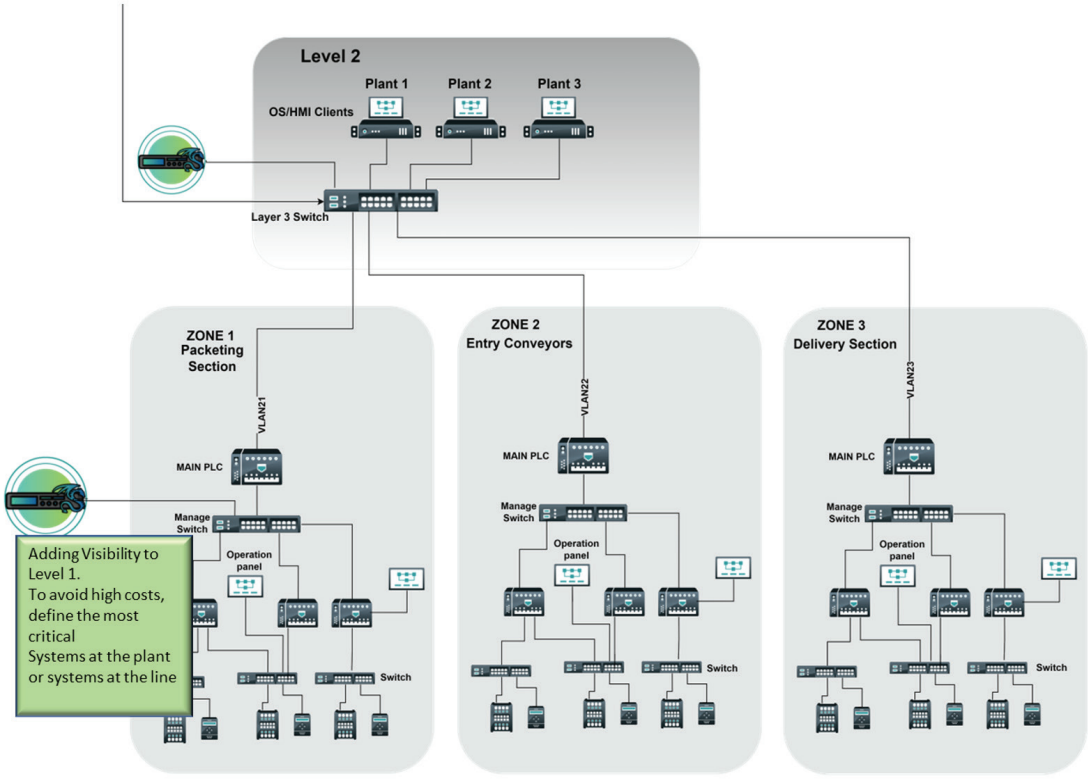

Figure 9: Network Diagram Adding Visibility for Critical Systems

WHAT IS A CROWN JEWEL ANALYSIS?

A Crown Jewel Analysis (CJA) is an iterative process that works top-down to systematically determine the physical and logical assets, data, and communication and control interfaces required for primary system function. Knowing the specific devices required for operation enables every aspect of vulnerability management, incident response, disaster recovery, and where protection and detection should be prioritized.

A Crown Jewel Analysis is comprised of six layers that define critical systems and their components. Layer 5 defines the controllers which handle the operation and monitoring of the defined critical systems components, while Layer 6 defines the crown jewels. Specific crown jewels can consist of multiple types of components including critical assets, communication paths and any functions which, if impacted, can affect the operation of critical systems.

Conclusion

Designing network architecture within the manufacturing sector presents challenges related to technology, connectivity, and equipment diversity. To address these challenges effectively, it is recommended to strategically separate each zone with systems and plants, as depicted in **Figure 10**. Collaboration between engineers, integrators, and cybersecurity professionals is crucial to establish necessary security controls and implement appropriate segmentation, starting from Level 1 of the Purdue Model.

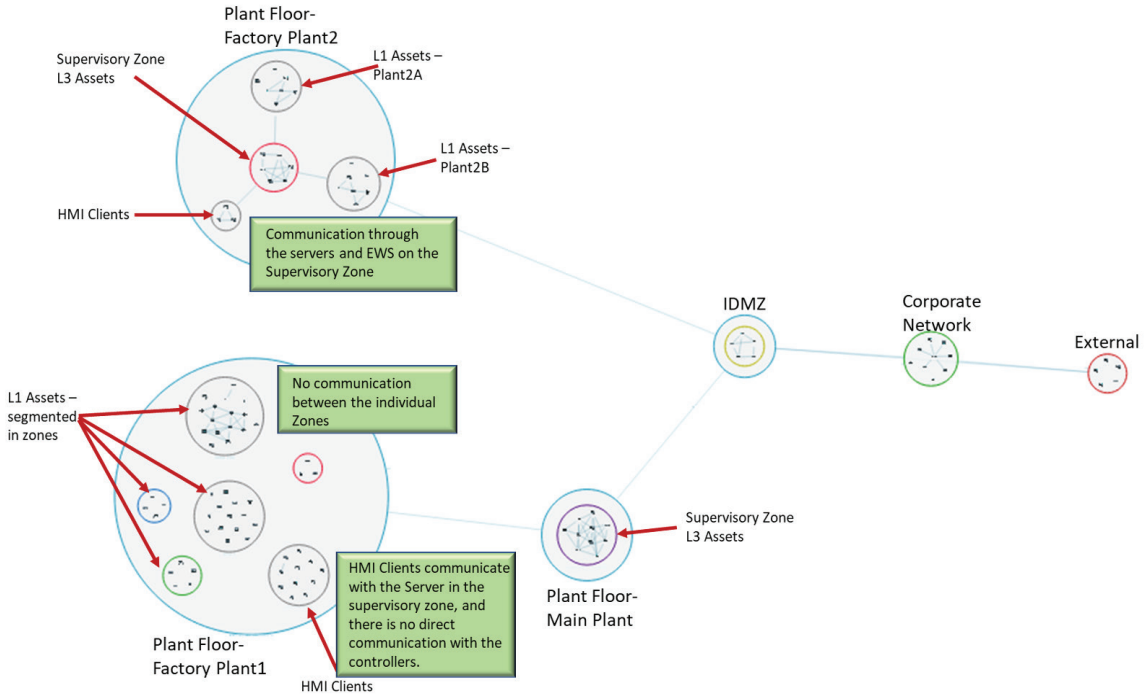

Figure 10: Dragos Platform Assets Map Level 1, Zones Segmented

Furthermore, the involvement of engineers is essential as these architectures facilitate the integration of automated tasks and enhance network performance. Centralized systems at Level 1 enable data loading for HMIs, MES, and analytics, establishing direct interaction with plant floor controllers. This centralization also streamlines the management of data and signal exchange between internal plants and processes. From a cybersecurity perspective, implementing segmentation and zone separation simplifies the identification of choke points, enabling effective monitoring to eliminate blind spots.

To improve monitoring capabilities, it is recommended to begin with Level 3 systems at the core switch (which is often shared with the DMZ) to monitor north-to-south traffic within the network. Subsequently, visibility should be extended to the switch that separates the interfaces between DB or HMI servers and the controllers, as well as the HMI clients. This approach enhances visibility for east and west traffic, concentrating on the interfaces between zones and servers. Once Level 1 zones are defined, visibility should be extended to all Level 1 systems, prioritizing critical zones based on Crown Jewel Analysis.

While the work of OT and ICS security professionals is never complete, striving to design an efficient and secure network architecture, as illustrated in **Figure 11**, can significantly enhance both ICS/OT security and automation capabilities in manufacturing sites. This approach provides a strategic advantage and robust defense for manufacturers, regardless of their specific production lines or overall business size.

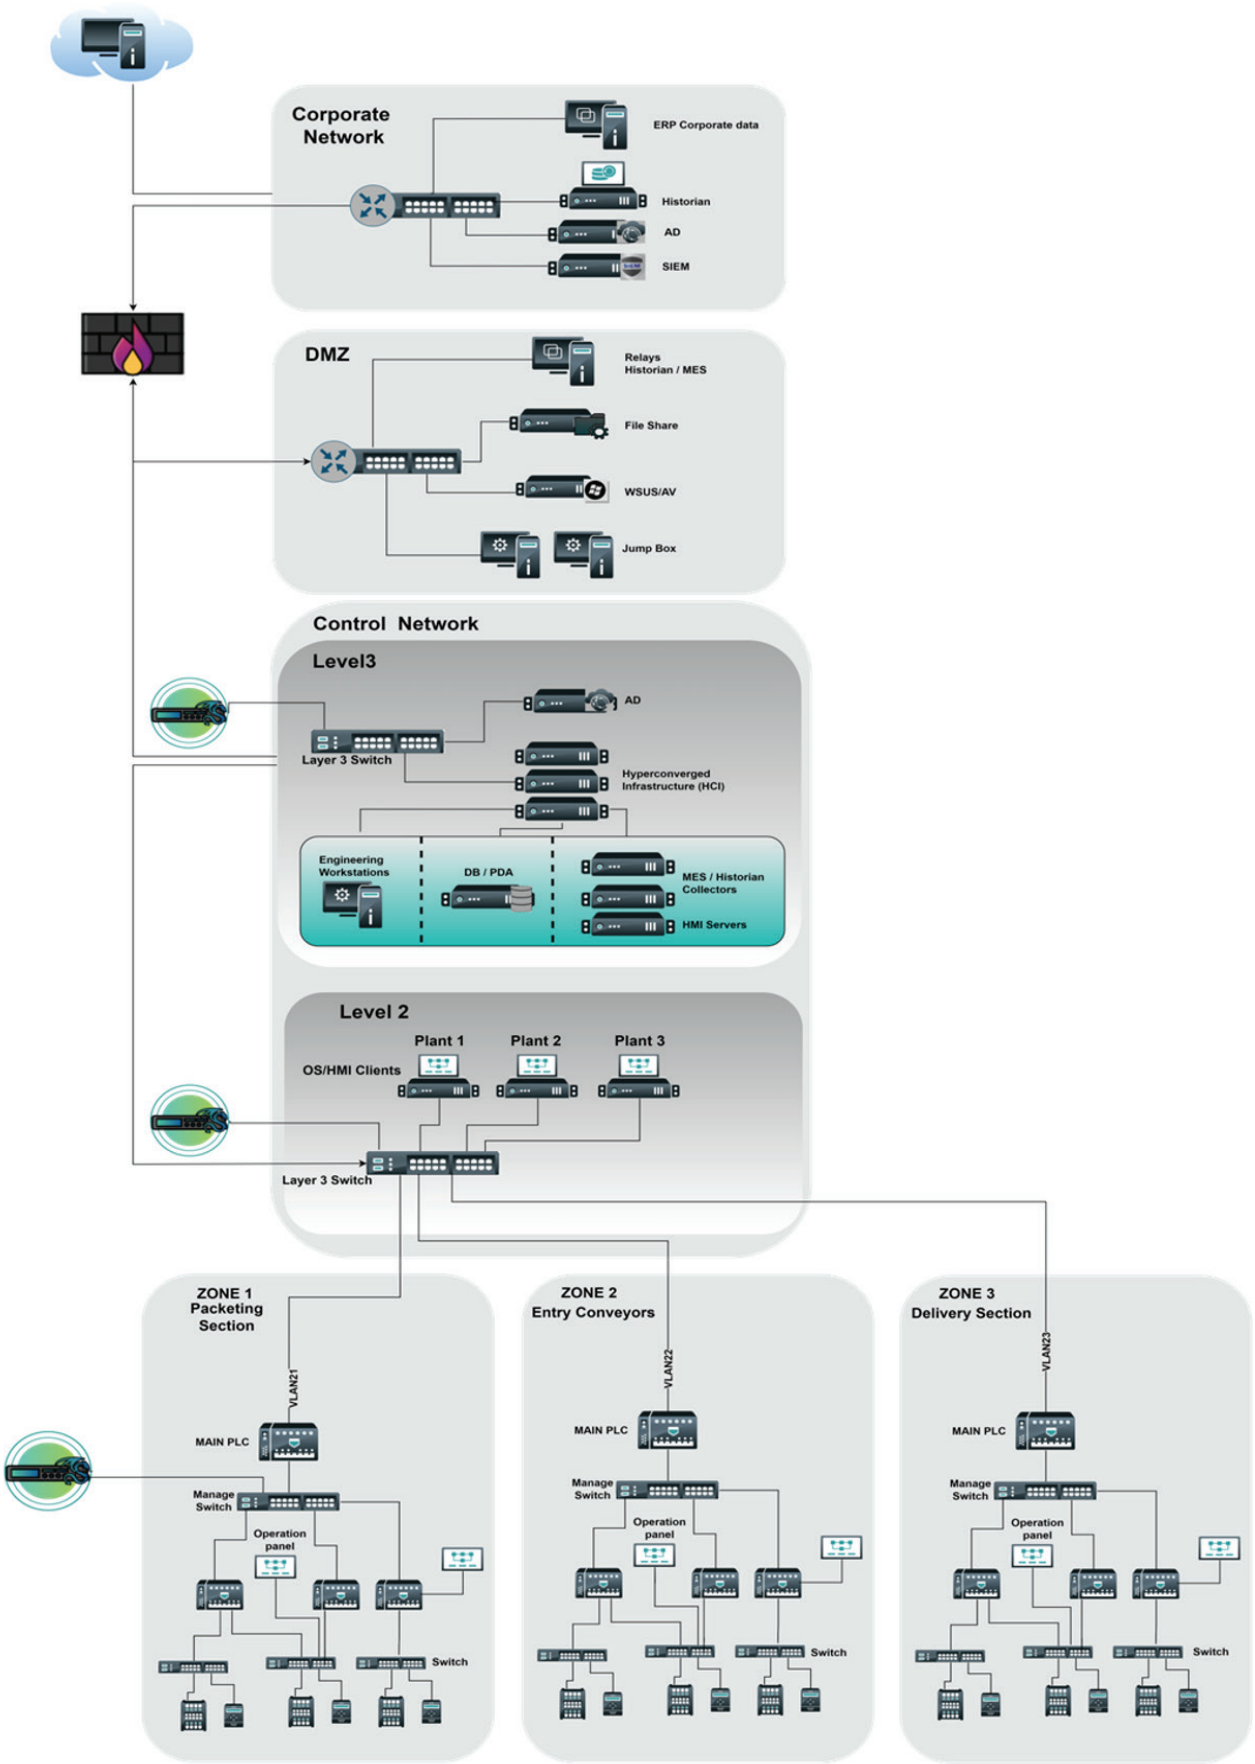

Figure 11: Final Network Architecture

Learn more about Dragos solutions for manufacturers:
Download our solutions brief.

About Dragos, Inc.

Dragos, Inc. has a global mission to safeguard civilization from those trying to disrupt the industrial infrastructure we depend on every day. Dragos is privately held and headquartered in the Washington, DC area with regional presence around the world, including Canada, Australia, New Zealand, Europe, and the Middle East.

To learn more about our technology, services, and threat intelligence offerings, visit **dragos.com** or connect with us at **sales@dragos.com**.