OT/IT Convergence by Data

In the past few years, the convergence of OT and IT (Operational Technology and Information Technology) has moved from curiosity to necessity in just about every industrial sector. Demands for improved efficiency brought on by the possibilities of Industrial IoT and Industrie 4.0 are compelling management to find better ways to extract value from OT data.

However, IT and OT systems are not particularly compatible. These two siblings have grown up separately, despite their common parentage in computing technology. They were designed for different tasks, with different goals in mind.

The focus of IT is business improvement—to support accounting, logistics, human resources, and all other areas of the business to make it more effective and productive. In a sense, for IT, *the product is the business itself*. Upgrades to computer systems and improvements in skills pay off with immediate results in the success of the business. And it's easy to make improvements because critical data is relatively static, providing ample opportunities to upgrade the tools and skills needed to manipulate the data.

In the OT world, the focus is on doing or making things. The *production process* is paramount. Complex factory systems, pipelines, power grids, and chemical plants cannot be switched on and off easily. Mission-critical systems running 24/7 cannot be put on pause for software upgrades. Every hour of lost production time can cost millions. It may take months or years to build such a system, and once it is running, few engineers are willing to risk swapping in a piece of untested software. Computer skills are just one aspect of a project where the bulk of the expenditure and expertise is focused on the machinery and devices needed to do the work. OT is one of several players in the game, and not the star of the show that IT often becomes in its world.

Convergence by Data

Given these fundamental differences, the best way for these two worlds to converge is by data. Specifically, production data from the OT side can be extracted and fed to tools and systems on the IT side. There it can power algorithms and analytics needed to enhance corporate performance. If necessary, data in the form of supervisory control commands, set points, or related information can also be sent back to the OT side. Taking this data-centric approach, OT and IT can effectively converge, while each remains in its own realm.

This approach depends on these essential elements:

- **Secure data access** is absolutely critical for both OT and IT, which means at least network segmentation, DMZs, and keeping all inbound firewall ports closed.
- **Protocol conversion** can enable data abstraction to a universal format.
- **Aggregation** eases congestion and simplifies data ingestion.
- **Edge processing** cuts bandwidth and reduces calculation loads.

- **Real-time performance** ensures the most up-to-date and accurate results.
- **Optional bi-directional data flow** can allow for supervisory control.

Secure data access - network segmentation

The most secure way to connect OT and IT, recommended by both industry and governmental agencies, is to segregate networks by using a DMZ ("demilitarized zone"). The NIS 2 Directive from the European Commission mandates higher security for networking data between the production and corporate levels of a company, encouraging the use of DMZs. On the other side of the pond, a White House memo advises industrial companies to separate corporate and production systems, limit access to OT networks, and ensure that they can continue to operate in the event that IT networks are compromised.

A reference document for both directives, NIST SP-800-82, sums up zero-trust OT network segmentation like this: "The most secure, manageable, and scalable control network and corporate network segregation architectures are typically based on a system with at least three zones, incorporating one or more DMZs."

These three zones are the control zone (OT), the corporate zone (IT), and the DMZ itself. Using a DMZ ensures that there is no direct link between corporate networks and control networks, and that only known and authenticated actors can enter the system at all. The SP-800-82 document describes the value and use of firewalls to separate these zones, and to ensure that only the correct data passes from one to the other.

Secure data access – implementing a DMZ

Properly implementing a DMZ between OT and IT networks is challenging—but possible. The challenge is that the two most common protocols used for accessing plant data in IoT systems, OPC UA and MQTT, were not designed for connecting via a DMZ. Getting data out of a plant through a DMZ typically requires two or more servers, chained together one after the other. The OPC UA protocol is simply too complex to reproduce well in a daisy chain like this. Information will be lost in the first hop. The synchronous multi-hop interactions needed to pass data across a DMZ would be fragile on all but the most reliable networks and would result in high latencies. And there would be no access to the data at each node in the chain. MQTT, on the other hand, can be chained but it requires each node in the chain to be aware that it is part of the chain, and to be individually configured. The QoS guarantees in MQTT cannot propagate through the chain, making data at the ends of the chain unreliable.

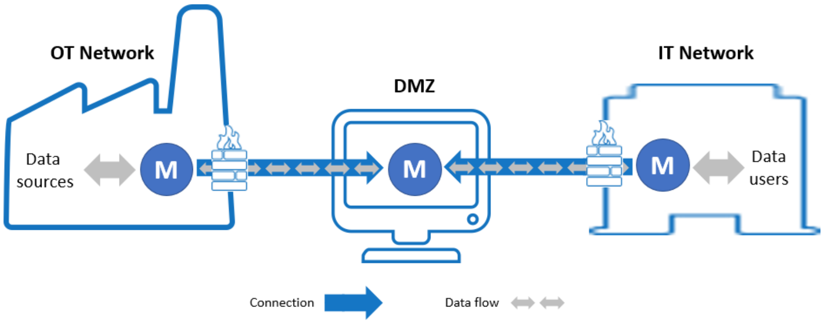

What is needed is secure tunnelling. Put simply, tunnelling means encapsulating data, allowing it to move from one network to another. Secure tunnelling middleware can support daisy-chained servers across a DMZ as long as it can mirror the full data set at each node. Implemented properly, secure tunnelling can provide access to that data both for local clients, as well as to the next node in the chain. The tunnelling middleware used should be able to guarantee consistency, so that any client or intermediate point in the chain will be consistent with the original source. Using a sophisticated tunneller in the DMZ can bridge between inbound connections from both OT and IT, ensuring that both firewalls remain closed.

Secure data access – closing all inbound firewall ports

To provide the highest level of cybersecurity, secure tunnelling must keep all inbound firewall ports closed. This is something that most industrial protocols were not designed to do. For example, OPC DA and OPC UA both use a client/server architecture, in which the client initiates a connection, and the server accepts it. The server must listen for the connection on a TCP port, and that port must be open for incoming connections on the server's firewall, and on any other upstream firewalls between the client and the server. To provide access to the data via OPC means opening at least one port on each of those firewalls, a significant risk.

Any open incoming firewall port constitutes a security exposure. Network attacks are not made on a port, nor are they made on a protocol – they are made on an application. The risk is that the listening application has an exploitable flaw that may or may not be due to the protocol implementation.

For example, an application may perfectly implement the OPC UA protocol, yet be vulnerable to flaws in OpenSSL. The application may have no exploitable flaws in OPC UA or SSL, but still have a buffer overrun in a string processing function executed on incoming data. No application is free of bugs. Every open incoming port is an opportunity for an attacker to probe an application for exploitable flaws and can give instant access to a network if one is discovered.

The best security practice is to not open any incoming ports at all in the secure network's firewall. With the right design, a tunnelling approach can meet this requirement. If the tunnelling middleware makes an outbound TCP connection from the server side to the client side, then there is no need to open any inbound firewall ports. This eliminates the attack surface altogether.

Protocol conversion

Multiple data protocols are common on both sides of a typical OT/IT data connection. On the OT side, the wide variety of serial, fieldbus, and other real-time data protocols are typically available as OPC DA or OPC UA. IoT data is often in MQTT format. Alarm data may be available as OPC A&E or OPC UA A&C, while historical data could be in OPC HDA, OPC UA, or ODBC. On the IT side, data is most often accessed through a dedicated data historian, an SQL database that uses ODBC, or Excel spreadsheets that use DDE. And for both OT and IT, proprietary, product-specific, or independently developed data formats may be implemented as well.

A data-centric approach to OT/IT convergence requires the ability to convert between these different protocols. Conveniently, some OT and IT products offer multiple data output or input formats, which allow direct connections between them. Should this not be available, or if other protocols must be included, a middleware solution can be used for protocol conversion. A well-designed solution of this type could convert all incoming and outgoing data streams to a single, universal format to keep the system as simple as possible, and automatically provide interconnectivity between any two connecting clients.

Data aggregation

Data on the OT side may come from a single source, but often it is necessary to collect it from multiple sources. IoT data streams coming from devices in the field can number in the hundreds or even thousands. It can be convenient or cost-effective to aggregate these sources into a single stream of data. For example, a cloud service may only allow a single data connection or might charge an extra fee for multiple connections.

Edge processing

Edge processing in this context means bringing computing power closer to the data source. For OT/IT data connections, edge processing offers these advantages:

- **Volume**: Industrial systems churn out enormous volumes of data, most of which is irrelevant. Edge computing can monitor the data and filter out what is unnecessary. This reduces bandwidth and frees up centralized or cloud-computing resources.
- **Cost**: Related to volume, feeding large quantities of raw data to an IT system or the cloud for processing is not cost effective. It is more economical to at least filter the data, or better still, process it locally and send the relevant results to the cloud.
- **Immediacy**: For any mission-critical system, the closer you can get to real-time decision-making, the better. Running right on the device itself, an edge-processing system can respond in a few milliseconds, compared to a round-trip to a central server or cloud system, which would take at least 100 milliseconds, and often longer.

Real-time performance

Although not absolutely essential, real-time performance is a big plus for OT/IT convergence. Real-time data coming directly from the factory floor gives the system immediacy, accuracy, and relevance not available through MES or ERP systems. Using real-time data to power analytical engines and predictive technologies keeps the company agile, able to respond more quickly to changing conditions.

Optional bi-directional data flow

For many users, true convergence of OT and IT means applying the analytical power gained from the IT side back into OT systems. This can come in various forms, such as closed loop supervisory control, predictive maintenance, digital twinning, and more. Since OT data is being used to build efficiencies in other parts of the enterprise, why not in the industrial process itself?

Any OT/IT data solution that meets the criteria we have discussed would be an excellent candidate for bi-directional data flow support. Secure transmission between isolated networks over one or more DMZs, through closed firewall ports, is equally essential for IT-to-OT data flow. Data would most likely need to be converted from IT protocols back into OT protocols. The system should also be able to redistribute data back to individual elements that it aggregates from. We would expect a bi-directional system to interface smoothly between IT systems and corresponding edge computing systems on the OT side. And to gain the most benefit from bi-directional data flow, real-time performance would be valuable.

Convergence becomes reality

Despite decades of skepticism that OT and IT might never converge, in just a few years this odd-couple match has become a reality. The focus of each world has not changed that much. IT is still concerned with business improvement, and OT continues to focus on doing and making things. But convergence by data has enhanced both worlds. Taking data from OT allows IT to bring the enterprise to even higher levels of efficiency and cost savings. And supervisory control based on the analytical tools of IT allows OT to cut costs and enhance productivity.