Introduction to HART-IP Security

FieldComm Group™
Connecting the World of Process Automation

Release Date: March 15, 2023
Document Number: FCG_AG10197

Document Distribution / Maintenance Control / Document Approval

To obtain information concerning document distribution control, maintenance control and document approval please contact FieldComm Group at the address shown below.

Trademark Information

Attention: FieldComm Group, President
FieldComm Group
9430 Research Boulevard Suite I-120
Austin, TX 78759, USA Tel: +1 (512) 7792-2300

http://www.fieldcommgroup.org

Use of imperatives in HART Specifications

The key words (imperatives) "must", "required", "shall", "should", "recommended", "may", and "optional" when used in this document are to be interpreted as follows:

Must	Must, Shall, or Required denotes an absolute mandatory requirement. For example, "All HART Field Devices must implement all Universal Commands"
Should	Should or Recommended indicates a requirement that, given good cause/reason, can be ignored. However, the consequences of ignoring the requirement must be fully understood and well justified before doing so.
May	May or Optional identifies a requirement that is completely optional and can be supported at the discretion of the implementation. May can be used to identify optional Host Application or Master functionality and, when this is the case, does not imply the function is optional in Field Devices.

Intellectual Property Rights

The FieldComm Group (the Group) does not knowingly use or incorporate any information or data into the HART, FOUNDATION Fieldbus and FDI protocol standards, which the Group does not own or have lawful rights to use. Should the Group receive any notification regarding the existence of any conflicting private IPR, the Group will review the disclosure and either (A) determine there is no conflict; (B) resolve the conflict with the IPR owner; or (C) modify the standard to remove the conflicting requirement. In no case does the Group encourage implementers to infringe on any individual's or organization's IPR.

TABLE OF CONTENTS

1 TARGET AUDIENCE

This is a high-level introduction to the mandatory security requirements provided by HART-IP servers/devices. It will benefit both users evaluating HART-IP or planning their deployment. It is also a good place to start if you are a device designer or thinking about developing a HART-IP product.

2 EXECUTIVE SUMMARY

HART-IP employs a defense-in-depth security strategy that blends prevention, detection, and deterrence. As a result, HART-IP is relatively simple to deploy and an effective tool that can enhance the plant's process automation security. This introduction begins with an overview of the HART-IP security triangle; communication, diagnostics, and forensics. To provide context, a brief overview of HART-IP is included.

Since secure HART-IP operation is a tool, its place in the plant security lifecycle is also reviewed. Finally, as many HART 4-20mA developers are not IP or security experts, integration of security into the product development lifecycle is summarized.

Secure HART-IP operation includes adopting the HART-IP security specifications, development of an overall secure device, and integration into plant operations and lifecycles.

3 INTRODUCTION

HART-IP, first deployed in 2009, provides a high bandwidth, easy to connect and use communication channel between Client applications and HART-enabled Servers (e.g., field devices and I/O systems). When used with I/O Systems, HART-IP is an open, interoperable backbone to access both WirelessHART and HART 4-20mA process instrumentation. HART-IP enabled field devices provide 20x faster process value updates, configuration uploads and diagnostic access (e.g., valve signatures, radar level waveforms) than traditional HART 4-20mA. Since 2009, HART-IP has been deployed in tens of thousands of networks (principally in I/O Systems).

HART-IP servers support at least 5 Clients which can enable applications like Monitoring + Optimization, SCADA, Control, PLCs, and Plant Asset Management (PAM) systems. Because HART-IP is IP-enabled, it is Physical Layer independent and works over Ethernet, Wi-Fi, satellites, cellular networks, packet radio, etc. Anything supporting IP communication. In process operations, including hazardous locations, HART-IP operates seamlessly over the Ethernet Advanced Physical Layer (Ethernet-APL) using twisted-pair wiring (10BASE-T1L).

Traditionally, HART Communication has been local in nature. HART 4-20mA is principally point-to-point over the current loop. WirelessHART, a mesh network that spans a process unit, still targets intra-plant communication. However, with HART-IP, communication can be global in reach. Consequently, robust cybersecurity is paramount. HART-IP security has always been mandatory. While the initial security choices were left to the vendor, since mid-2020 the minimum-security requirements have been standardized. This paper introduces HART-IP's standard security arsenal.

HART-IP embraces a three-prong, defense-in-depth security strategy (see Figure 1). The first is security on the HART-IP communication itself. This is a prevention-oriented strategy and is built upon industry-standard Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS). TLS/DTLS is the foundation for banking, e-commerce, email, and all the other applications you already use over the internet.

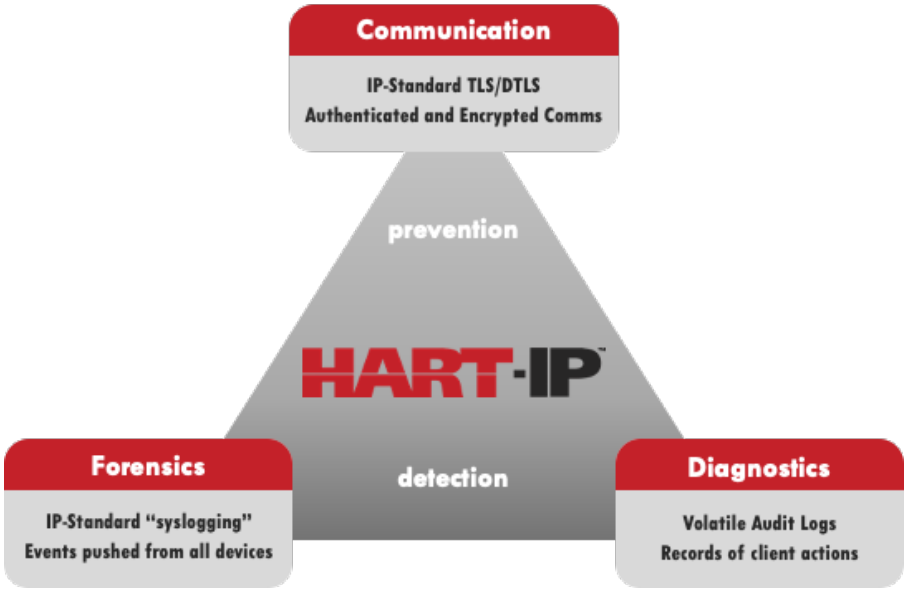

Figure 1 – HART-IP Security Triangle

Communications security is augmented by two detection-oriented strategies: diagnostic Audit Logs and syslogging (forensics). syslogging is also a widely used industry standard to provide network wide troubleshooting and forensics. syslogging allows events to be recorded (with timestamps) from all network devices (not just HART-IP devices). These can be analyzed in case of a security breach, unscheduled plant shutdowns and other plant wide events. Audit Logs are event and usage records stored in volatile memory that provide diagnostic information. For example, Audit Logs record what clients have communicated with the device and even which client changed the device's configuration.

Prevention and detection are complementary. Prevention (communication security) provides robust protection against incursions. Detection (syslogging and Audit Logs) enables the identification of the events and agents that penetrated or disrupted operations. Detection offers an extra layer of protection by providing a strong deterrent.

HART-IP security was designed using the process industry expertise of the FieldComm Group and its member companies. It is effective, robust, and scalable in small air-gapped installations and large enterprise-wide applications. A minimalist installation is achieved by deploying HART-IP security TLS/DTLS plus Audit Logs with little to no other network infrastructure required. In larger installations, assets like syslogging, and dedicated security managers can be deployed by plant network engineers.

4 OVERVIEW OF HART-IP

HART-IP combines the ubiquitous Internet Protocol (IP) and the HART network and application layers. Consequently, plant and corporate personnel can utilize infrastructure already deployed and understood to provide HART compatible system connectivity and integration. HART-IP is not limited to only Ethernet but works over any communication medium (Physical Layer) that supports the Internet Protocol.

The following subsections provide an overview of the HART-IP architecture and introduces core HART-IP features including HART-IP Devices, the HART Application Layer, the layering of HART-IP on top of the ubiquitous

Internet Protocol (IP) and how HART-IP complies with HART Protocol Requirements for interoperability and backward compatibility.

4.1 HART-IP DEVICES

HART-IP is a Client/Server protocol. Servers are, for example, a process instrument or an I/O System. An I/O system provides connection(s) to HART 4-20mA or WirelessHART process instruments. Clients are HART-IP enabled host applications like Monitoring + Optimization, SCADA, Control, PLCs, and Plant Asset Management (PAM) Applications.

Figure 2 depicts a network connecting HART-IP Devices (Clients and Servers). Clients are arrayed across the top of the figure including a Historian, Monitoring and Optimization (M+O), PAM, and a Controller. There is a firewall protecting the instrumentation / control network. Servers are across the bottom. On the left is a WirelessHART Gateway providing Client's access to all the process instruments on that sub-network. Similarly, Remote I/O (shown on the right) provides access to HART 4-20mA instrument's runtime process and status data, configuration, etc. In the center native HART-IP process instruments are shown.

HART-IP is session-oriented. In other words, Clients must open a session before communicating with the HART-IP Server. Clients can communicate with the HART-IP device continuously or intermittently. Historians and Controllers will generally keep a session open indefinitely. PAMs are usually task-oriented and thus intermittently connect to the Server as needed. When the client is finished interacting with the server it should close the session. Servers support multiple sessions to allow simultaneous access by several clients.

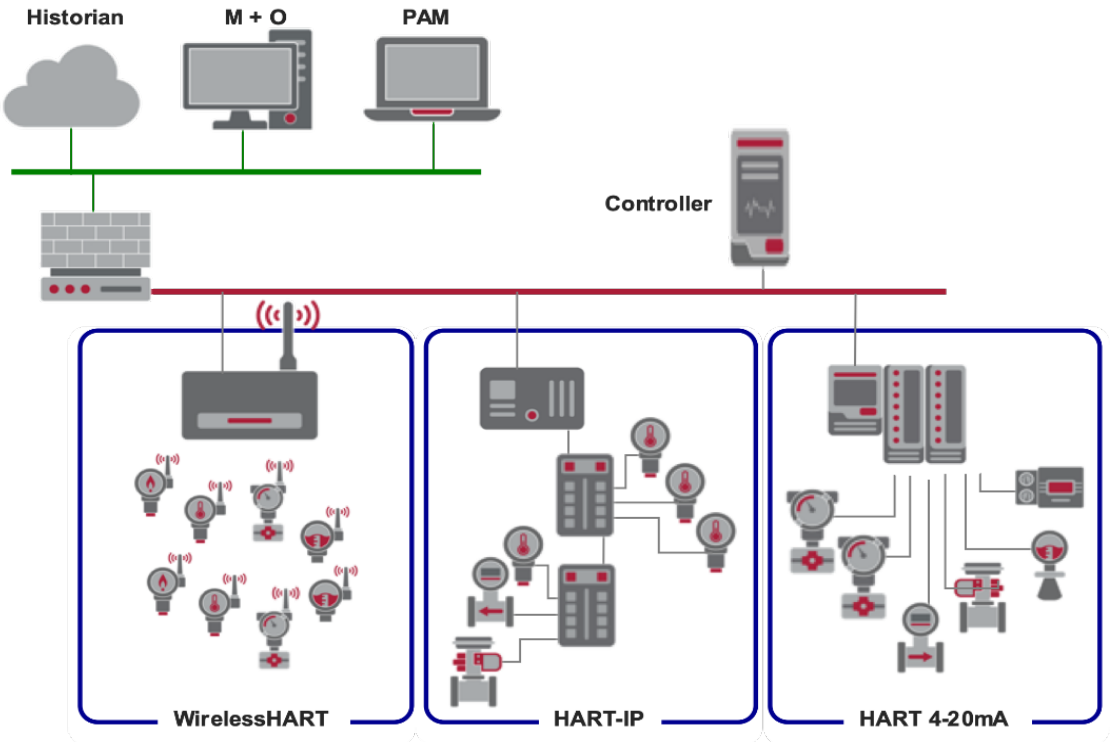

Figure 2. HART-IP Devices and Communications

In addition to asynchronous request/response communications, HART-IP Servers can publish runtime process and status data. Generally, the process engineer configures the publishing criteria based on the connected process. Clients simply subscribe to the publish stream from the process instrument and the data will be pushed to them. This is very efficient allowing Servers to push process and status data to multiple clients without repetitive polling by the Client. Publish / Subscribe is especially beneficial when the data is pushed from a WirelessHART or HART 4-20mA subnetwork where bandwidth is more limited. In this case the I/O System will receive the published data and distribute to all the Client subscribers.

4.2 HART Common Application Layer

All 60M+ HART-IP, WirelessHART, and HART 4-20mA devices support the same Application Layer. That means all HART servers (Field Devices and I/O) have standardized identification including Manufacturer Name, Device Type, Device Revision, and Tag. In addition, every HART device has a Unique Identifier (ID) that is fixed and unchanging. HART-enabled PAM systems use the 40-bit Unique ID to track HART devices through their lifecycle. Likewise, access to runtime process (process values with status) and status data is standardized. The HART application layer encompasses over a half-dozen Specifications including those for Universal and Common Practice Commands.

> **What are HART Commands?** Commands are the methodology for encapsulating access to HART device functions and data. In the Protocol, Commands are cohesively grouped to access device features/functions. There are groups of commands to access process values; manage the loop current; calibrate sensors; manage publisher/subscriber communication; access WirelessHART or HART 4-20mA devices via a HART I/O; and so forth. Commands are the widow into the device's application.

HART-IP field devices provide the same capabilities as found in WirelessHART or HART 4-20mA field devices. Command 0 provides the basic identity information (device type, device profile, manufacturer, device revision, etc.) The device's Tag is read using Command 20; process data via Commands 1-3, 9; Command 48 reads extended status information; and Command 38 the Configuration Changed Counter. The same data and commands as found in WirelessHART and HART 4-20mA.

WirelessHART Gateways and Remote I/O have supported HART-IP since 2009 via standard I/O System Commands. That means there is a standardized way to talk to instruments connected to the I/O.

Procedures and tooling are the same in all cases as well. For example, most HART PAM tools already support HART-IP. Aside from (like WirelessHART) there not being a current loop, plant staff can manage and use HART-IP like they do HART 4-20mA. Of course, HART-IP communication can be a lot faster: more process values per second; faster configuration uploads, etc.

4.3 HART-IP Layering

The technical approach to HART-IP is like that taken when MODBUS-TCP was developed - take the communication protocol that is already widely used, well known, and user-accepted then layer it on top of IP. Figure 3 shows the layering of HART-IP. The lower layers are the standard IP communications (including security).

4.3.1 The HART-IP Layers

The HART-IP layers add a simple HART-IP header that encapsulates a standard HART 4-20mA packet. All HART Applications understand this packet structure and the HART Application Layer. This makes it simple to support HART-IP because:

- HART-IP is built using the well-known HART principles
- IP stacks (including TLS/DTLS security) are readily available (e.g., open source, or commercial)
- Standard PHY and MAC (e.g., Ethernet, Wi-Fi, etc.) are built into many processor chips

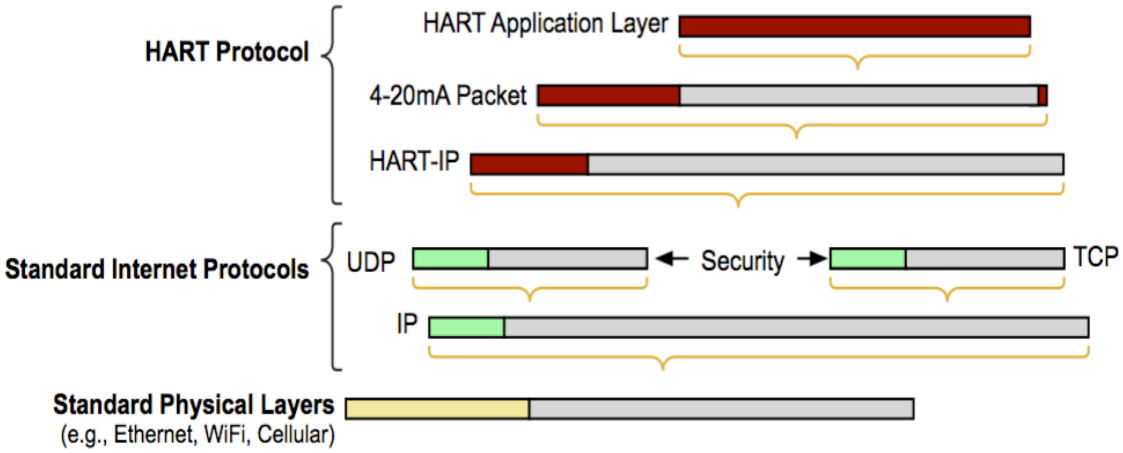

Figure 3. HART-IP Communications Layers

In other words, existing HART-aware host applications have found it relatively simple to adopt HART-IP. The host uses operating system services to connect to the HART-IP Server via IP. Support for the HART-IP PDU must be implemented. Since the Application Layer PDU is unchanged (the 4-20mA packet and Application Layer data), it can be extracted from the HART-IP datagram and the rest of the host application processing is unchanged. In fact, many host applications have transparently supported HART-IP enabled I/O systems to communicate with HART 4-20mA and WirelessHART field devices for many years.

Feedback from host vendors already supporting HART digital process values and status indicate rapid success in supporting HART-IP field devices. Simply add a driver to access the HART 4-20mA packet from the HART-IP stack and the remaining HART Application Layer data permeates the host application without any further modifications needed.

4.3.2 Network Services - The Internet Protocol Layers

When designing HART-IP the engineers in the HART Technology Working Group sought a careful balance between process automation practices, HART simplicity, and leveraging IP technology. The result is a specification that blends both HART and IP technology.

Standard network services supported by HART-IP and used in IP networks include:

- DHCP (Dynamic Host Configuration Protocol) used to assign IP addresses
- DNS (Domain Name System) that associates a HOSTNAME with an IP address

- NTP (Network Time Protocol) / PTP (Precision Time Protocol)
- Syslogging (see Subsection 0)

DNS and DHCP can be challenging and are well suited/familiar to IT technologists but less so to process automation users. Consequently, HART-IP leveraged DHCP and DNS while simplifying and adding a paradigm well-suited to process automation.

Process automation documentation, user consoles regularly use and think in instrument tags. HART devices have supported tags for decades and so, with HART-IP the TAG became the HOSTNAME. That means if you want to connect to a HART-IP device you simply connect to the TAG (which is the HOSTNAME in DNS). Since sometimes multiple process units may be identical all the way down to the tags, HART 7 added Process Unit Tags in 2007. Ambiguity can be reduced by entering the Process Unit Tag. Then, the HOSTNAME will become the concatenation of the TAG and the Process Unit Tag.

How to commission an instrument fresh out of the box was also considered. The factory default TAG (HOSTNAME) is the hyphenated hexadecimal MAC address (the physical address of the device). HART-IP requires manufacturers to physically print the MAC address on the outside of the device. That means if the device has not been provisioned yet you can still look up the device in DNS.

HART-IP also includes specifications that allow a simple, bench-top network **with no network services at all**. This allows the initial provisioning (see Subsection 5.1.2) of the HART-IP device to be performed in the instrument shop the way HART 4-20mA and WirelessHART devices are often done today.

Time synchronization is essential in process automation and, consequently, HART-IP stipulates support for industry standards NTP and (optionally) PTP.

As can be seen, HART-IP's integration with IP is thoughtful and includes paradigms helpful to process operations.

4.4 Interoperability and Backward Compatibility

Interoperability and Backward Compatibility are two of HART's central pillars. The HART protocol specifications define product requirements in great detail to ensure interoperability between products. It specifies requirements of how field devices and the protocol are to be improved. Fundamentally, new features can be added to a field device or the protocol, provided the existing features remain supported. It is a mandatory requirement that if an old field device is replaced with a new one, it should be compatible with the existing host applications. FieldComm Group's HART Registration program validates compliance with these requirements.

There are two versions of the HART-IP requirements, and, like field devices, Version 2 (v2) is compatible with Version 1 (v1). A v1 client can interoperate with a v2 Server. The FieldComm Group QA/QC program tests and confirms this requirement every time a HART-IP Server is registered.

With HART-IP, the compatibility is determined while the Client and Server open the session. That is before any HART communication even occurs. The Client tells the Server what version its supports. If the Client is v1, then the Server answers as v1. If the Client is v2 and the Server is V1, then the Client must tailor its communication accordingly. For example, the v2 Client knows it cannot utilize Audit Logs in v1 Servers (they didn't exist in v1).

Communication security is also a critical element of backward compatibility requirements. If both the Client and Server are v2, then security credentials are exchanged, and the communication channel is secured during the

session-initiate. The security negotiations comply with the TLS/DTLS standards. If the Client and Server credentials cannot be authenticated, then the session-initiate is aborted.

5 HART-IP SECURITY

The engineers designing HART-IP have extensive experience in the process industry and with their plant personnel. Their insights were leveraged while standardizing the minimum HART-IP security requirements. The key objectives for success include:

- Robust, industry-standard security
- Relatively simple to provision with strong ties to existing HART best practices
- Deployable on air-gapped networks (no internet access required)
- Scalable from small, isolated networks to large plant installations

The result is a three-legged, defense-in-depth strategy with one leg focused on prevention (communications security) and 2 legs focused on detection/deterrent (Audit Logs and syslogging). The following are the minimum features required in HART-IP products. (Manufacturers and end users are encouraged to augment these as needed to ensure secure and safe plant operations. In addition, FieldComm Group and its members recognize that security best practices for IP-based products continue to evolve. Consequently, all products must support field upgrades allowing the addition of improved security mechanisms in the future.)

5.1 COMMUNICATION SECURITY

Communication security is based upon TLS/DTLS versions 1.2 or later. TLS/DTLS supports about 40 cipher suites. Over time, new cipher suites are added and, when vulnerabilities are discovered, cipher suites may be retired. As the threat environment evolves, HART-IP Communication security will also evolve and improve. Thus, the requirement that all HART-IP devices be field-upgradable.

HART-IP currently requires Servers to support 4 cipher suites. These suites are both resilient to attack and can be implemented efficiently on simple process instruments with limited resources.

5.1.1 Overview

The operation of TLS/DTLS communication security consists broadly of two phases: authentication and operation. Authentication is the process of the client and server learning if they can trust each other. For this phase, HART-IP requires either Pre-Shared Key (PSK) or Secure Remote Password (SRP) authentication. Both are widely used in industry, especially in low-power embedded devices and air-gapped networks. While both must be supported in HART-IP Servers, the end-user can choose either or both for use in their process instrumentation network.

Once authentication is successful, AES-128 (like used in WirelessHART) is used to encrypt and decrypt individual communications. The 128-bit key used during operation is generated and exchanged securely during the authentication step. AES-128 is widely used in IP networks and has natively supported computation modules in many low power embedded microcontrollers.

User Credentials

Both PSK and SRP require security credentials to be provisioned in advance. HART-IP servers must support security credentials for at least 5 "users". A user may be a physical person or a metaphor for an actor like a security manager, data historian, plant operator, instrument technician, etc. Each user has an ID and either a PSK or Password (or both). Multiple client sessions may be opened using the same user credentials. For example, a system may have multiple operator consoles using the same User ID to subscribe to runtime process and status data.

Security Manager

While the credentials can be managed manually, HART-IP recommends the network have a dedicated Security Manager application (like WirelessHART does). Using a Security Manager ensures the credentials (PSK and Password) are confidential and sufficiently random to not be guessed easily (if at all). The "Security Manager" corresponds to the first (slot 0) set of user credentials. Only the Security Manager can change any of the user credentials.

Read-only

HART-IP allows one or more users to be designated "read-only". This user should subscribe to the runtime process and status data generated by the HART-IP Server. While read-only users can have request/response communication with the server, no transactions that can affect the configuration or operation of the Server are allowed. The offending client is disconnected in this case (the response is not even transmitted by the server).

5.1.2 Initial Provisioning

HART-IP envisions initial Server provisioning to be like that used with other HART devices (e.g., in the instrument shop on the bench). To perform this initial provisioning, all HART-IP servers must support operation on a minimal network as shown in Figure 4. No IP infrastructure (like DHCP, DNS, etc.) is required. In this case, the Server will default to IP address **192.168.0.42**. This allows the HART-IP Client to connect to the new device without any IP infrastructure as well.

The very first connection to a fresh HART-IP server or one that has just been Factory Reset is a critical junction in its life cycle. After this, the Server will begin its life in v1 backward compatible mode (no security) or in secure v2 mode with the credentials for the security manager established.

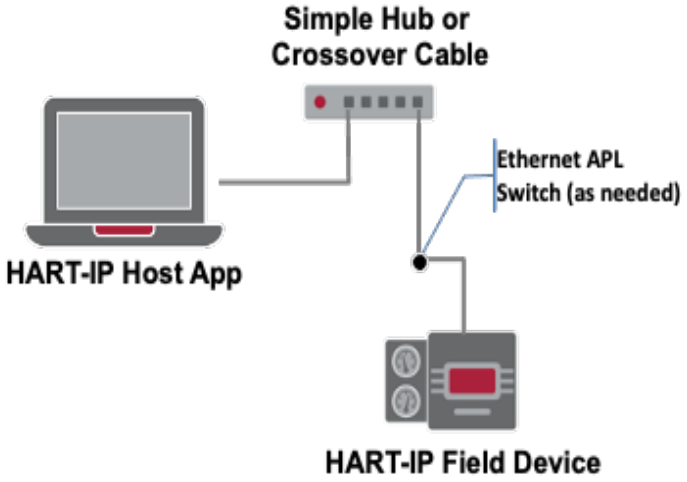

Figure 4. "Benchtop" HART-IP Network

Any Client (v1 or v2) will connect to the server during this initial session. If the initial connection is via a v1 Client then the Server is locked into Version 1 and no security credential can be entered. Future connections from v2 clients will be refused. Communication security must be provided by the end user (e.g., via a VPN).

If the initial session is with a v2 Client then Server is locked into Version 2 and secure communications must be initiated using the default authentication. The Client must perform the initial security provisioning during this initial session. The following properties should be provisioned:

- The Tag
- At least one set of security credentials (PSK or Password) for Slot 0 (the Security Manager)
- (optional) Process Unit Tag
- (optional) syslogging parameters

If a v2 Client does not enter security credentials during the initial session, then they are changed to cryptographically secure pseudo-random values. Communication is not possible with the device until a Factory Reset is performed.

The user should perform this initial provisioning on the bench along with all other routine device provisioning (e.g., setting unit codes, burst mode parameters, etc.). Performing the provisioning on an isolated network allows initial security credentials to be established unobserved.

Once the initial provisioning is complete the device can be deployed and the instrument's configuration completed. The device may remain quarantined for a while before it is made operational.

5.1.3 Factory Reset

All HART-IP servers must include manual/physical means to reset the device to its factory settings. This provides a secure means to recover the server in case, for example, its security credentials are lost; the device is decommissioned; or the network is being upgraded from v1 to v2 operation. When Factory Reset is performed

- All parameters, including calibration, etc., are reset to factory default values
- Slot 0 user credentials are set to their default values
- Tag (HOSTNAME) shall be initialized to the MAC address of the Server (required to be written and visible on the device itself) and the Process Unit Tag is cleared
- Supplemental UDP and TCP IP-ports shall be reset to default (5094)

After Factory Reset the Initial Provisioning must be performed.

5.2 AUDIT LOGS

The Audit Log summarizes Client activities. Basically, it records, on a session (connection) by session basis, which Client (and when) communicated with the Server and what the Client did during that connection. The Audit Logs include a circular list of at least 128 records of Client sessions. Each client session summary record consists of:

- Who. Client identity (IP address and port number)

- When. Connect/disconnect time/date
- Starting/ending Configuration Change Counter. Was the configuration changed while that client was connected?
- Communications summary for that session. Counters indicating the number of publish, request, and response PDUs during that session.

The Audit Logs provide Diagnostic summaries that can detect which clients and thus, who has been interacting with the server. It could be used, for example, to see who inadvertently changed the device's configuration.

A high-level summary of the Server operation itself is also provided. This Includes:

- Last time/date the server was powered up
- The last time/date the security credentials were modified
- Server Status (e.g., was the syslog server located)

Since the Audit Logs (while extensive) are volatile, it could be power cycled to clear them. However, the time when the Server is powered up can be used to detect this kind of mischief.

5.3 SYSLOGS

Unlike Audit Logs, syslogging requires network services (specifically a syslog server) to operate. However, while Audit Logs provide summaries, syslogging provides an event-by-event history of Client and Server activity. syslogging is a standard IT and security tool first introduced in the 1980s. syslog messages are pushed from all network devices (including HART-IP devices) and sent to the syslog server. The messages include time stamps that allow the sequence of plant-wide events to be ordered and sequenced.

HART-IP Server pushes a syslog message when, for example:

- The server starts up or shuts down
- Security credentials are changed
- Any configuration parameter is changed (and which one)
- A client session is opened or closed

syslog servers are databases ranging from very simple to sophisticated Security Information and Event Management (SIEM) systems and are critical to detecting security attacks and unauthorized manipulation of network devices. These servers answer queries that enable in-depth forensics to detect and unwind network-wide anomalies or mischief. It can also be used to detect possible attacks. For example, HART-IP servers push a syslog message when a client unsuccessfully tries to establish a connection. This might be a sign of an attacker probing the network.

6 PLANT SECURITY LIFECYCLE

HART-IP provides robust tools for ensuring secure operation. To be effective, end users should take a systems view of cyber security with HART-IP security tools being one part of the solution. Appropriate security management must also be implemented (for example in managing security keys, passwords, unauthorized access, etc.).

HART-IP devices are installed and will remain operational for 15 to 30 years. Over the device's lifespan, plant overhauls, process modifications, and business integration activities can have security implications. Consequently, End users should manage security through the device's entire life cycle. The phases in the device's security lifecycle are illustrated in Figure 5.

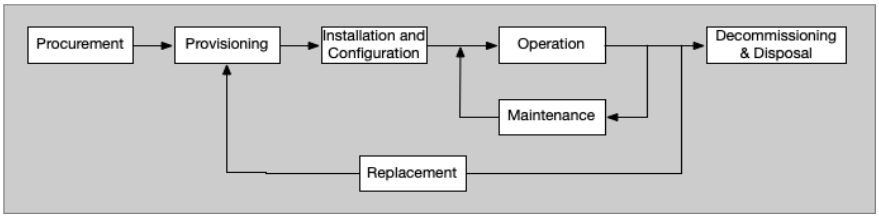

Figure 5. HART-IP Device Lifecycle

Example activities in each phase of the security lifecycle include:

Procurement Procurement is outside of the scope of this technical paper.

Provisioning Devices should be provisioned (onboarded) before installation (see Subsection 5.1.2). This includes basic provisioning (like writing the device's Tag) and writing security credentials. Initial provisioning should be performed in a secure location/network.

Installation and Configuration Following provisioning, the device is installed in the plant. Since security has been provisioned, it will be admitted to the plant network or the air-gapped subnetwork. The device should be configured for operation. This may be done manually by the PAM system operator, or the configuration may have been preloaded in the PAM. In the latter case, the configuration can be downloaded to the device.

To improve security robustness, all security credentials may be updated (changed) directly by a security manager during the installation process.

Once the installation is complete a baseline snapshot of the configuration should be captured and verified.

Operation Once the device (server) is installed and in production, its integrity should be monitored regularly. HART-IP enables the integrity to be evaluated at several levels. First, device publishing defaults to including device status and the configuration changed counter. A change in status may indicate a maintenance issue and any increment in the configuration change counter indicates the device's configuration has been manipulated. Both can be investigated using the PAM tools.

The Audit Log can be accessed and the source of any configuration change can be identified. Also, the Audit Log indicates any change in the security credentials.

Since syslog messages are produced regularly by HART-IP devices, many (SIEM) systems can be configured to automatically alert when key events (like configuration changes) occur. Similarly, the SEIM can be consulted should a plant-wide event (like an unscheduled shutdown) occur. Since all syslog messages are time-stamped the order of occurrence can be determined.

As a preemptive measure, security credentials in all HART-IP devices should be periodically updated.

Maintenance

Plants have established maintenance procedures/policies. Maintenance may be performed in-situ or, during a plant shutdown, moved to a maintenance depot. If the device is removed for maintenance while the plant is in operation the **Replacement** activity may be applicable.

In addition to typical device maintenance, a firmware upgrade may be required. All HART-IP Servers must support field firmware upgrades. Communication security is continuously evolving as threats emerge and improvements to the device's security suite may be necessary.

Prior to returning to service a snapshot of the device configuration should be made and verified.

Replacement

Occasionally a device must be replaced. Provided the device is being replaced by the same device type and the same or later device revision[1], the replacement of a HART device should go smoothly. If possible, a snapshot of the existing device's configuration should be taken before removal.

The replacement device should complete **Provisioning** prior to its installation. Security credentials (e.g., PSK and Passwords) should be unique to the replacement device. Following provisioning the replacement can be installed and the stored configuration downloaded to the replacement device.

PAM track HART devices through their lifecycle using the devices' Unique ID. Consequently, the replacement device may be a new record in PAM (a new Unique ID). A snapshot of the replacement device's configuration should be made and verified.

Decommissioning & Disposal

Whether the device was replaced and moved to maintenance stores or it is being discarded, security should also be evaluated. In either case, at a minimum, performing a Factory Reset will reset the device configuration and security credentials to their default values. This minimizes data leakage that could be compromising.

[1] The. HART Protocol has specific requirement governing device revisions to ensure forward and the FieldComm Group QA/QC program confirms compliance when device revisions are registered.

If the device will be repaired and moved to maintenance stores it maybe provisioned and readied for re-installation. Of course, it could be used for training or process development as well.

If the device is destined for disposal further physical disassembly or recycling may be appropriate. For example, non-volatile memory elements might warrant destruction (consult with the device manufacturer).

7 IMPLEMENTING HART-IP SECURITY

Security assessments, requirements and planning are another axis of requirements that must be included when designing modern process instrumentation. This axis should be included in product design along with the usual axis (precision and accuracy, safety, usability, speed, reliability, etc.)

In fact, industry initiatives insist that security be considered at every step in the product lifecycle. Like with Quality Assurance (QA) and Quality Control (QC) principles, an independent person or team should be responsible for ensuring the development meets security requirements and protections from concept to end-of-life. In any case, security reviews must be held regularly during product development beginning at the concept phase.

7.1 SECURITY DEVELOPMENT LIFECYCLE

Development teams should have a consistent repeatable process to ensure customer needs are met, reliable products are delivered, and predictable schedules are attained. No matter whether the process adopted is the traditional waterfall, agile, or other development models the key activities include (for example) Requirements, Design, Implementation, Verification, and Release. Security considerations must be included in each of the key activities. Security activities should include:

- Requirements: identify security requirements; risk assessment; and establish security gates
- Design: assess surface threats; threat modeling; and security design
- Implementation: static analysis; discard unsafe/insecure functions; and use trusted, secure tooling
- Verification: dynamic analysis, fuzz testing, and review of attack surfaces
- Release: establish incident response plan; final security review; and archive the release

However, the Security Development Lifecycle spans more than just the normal product development cycle. The Security Development Lifecycle begins when initial product concepts are formulated, runs through the normal product development and continues until devices are retired from the field. During the concept phase security training should be provided to the development team prior to project initiation. Security processes must be established, and responsibilities allocated to staff. Once released and products move to the field, they should be monitored for security breaches. Manufacturers and end users should develop a security incident response plan and, if a security breach occurs, the incident response plan must be executed. This may result in firmware updates to the installed base.

Security must be embedded into the company development fabric and is a new axis that sources additional consideration into product development and release. When pursuing security certification, the security development lifecycle will be assessed and be a prerequisite to certifying a device.

7.2 SECURE DEVICES

The security of the device's firmware - its resistance to tampering - is paramount. Historically this could be assured by using One-Time-Programmable (OTP) memory. However, HART-IP must be field upgradable to respond to future security threats and necessary improvements to existing or new cipher suites. Given that, the device must be designed to assure that its firmware, in total, is what the manufacturer intended.

This begins by establishing a "Root of Trust" (RoT) that forms the founding link in the chain of trust used to validate all firmware in the field device. This anchor to the secure device can be accomplished by, for example, booting from read-only code (e.g., by disabling the JTAG on production devices), an independent coprocessor that validates the main field device code, or storing and operating the boot code from a protected space. Chip/SoC suppliers offer a variety of technical approaches to assist in establishing this RoT.

Even the initial, protected firmware should be self-validated. This initial boot code, sometimes called a Trusted Computing Base (TCB), is generally a relatively small piece of software and hardware that exists to confirm the integrity of itself, provide the critical security infrastructure, and confirm the authenticity and integrity of the field device application. If the boot code does not authenticate then it is not executed. In some designs, the secure boot code may also include basic I/O services (used by the application) as well as services needed to update the field device application code in the field.

Once the RoT is established the other firmware modules can be authenticated each in turn until the field device is fully operational. The same mantra applies - if the firmware module does not authenticate then it is not executed. This incremental authentication ensures the chain of trust remains intact.

The device's design must also provide secure storage for the identification of the field device itself, its clients, and the keys and passwords used to secure communications. Secure data items can include both confidential and protected data items. Confidential data are secret and must not be externally accessible. Confidential data (like security credentials) should be stored in secure memory space and data items themselves should be obfuscated. Protected data (Like the device's Unique ID or MAC address) are not a secret but the value normally never changes. In general, physical security cannot be ensured and devices should be designed accordingly[2].

7.3 FIELD FIRMWARE UPGRADES

HART-IP requires that servers support field firmware upgrades. In general, firmware upgrades could be performed as follows:

- Means and methods (e.g., using HART Block Data Transfer) are required to download the update into the field device. These need to work even when the field device is connected to an isolated network.
- The device should have sufficient memory to contain both the update and the corresponding operational firmware module.

---

[2] For example, researchers have shown that in some WirelessHART implementations it is easy to locate and access the Join Key by dumping NV memory thus allowing unauthorized devices to join the mesh network.

- The update must be validated. This could be done in the same fashion as when validating firmware modules on startup.
- If validation fails, the field device must reject the update and recover gracefully (rollback to a previous version).
- Once validated, the updated firmware must be brought into operation. Developers should retain the previous firmware version in case a firmware rollback becomes required.

Clearly support for upgrades requires planning, design, and testing during product development. Also, additional resources (like memory) are needed when designing the field device hardware.

8 CONCLUSION

HART-IP security consists of both prevention and detection strategies. Communication security (prevention) is based upon industry-standard TLS/DTLS. Audit Logs and syslogging augment communication security to detect and deter mischief.

Figure 6 depicts a HART-IP monitoring network with HART-IP field devices being Ethernet-APL enabled. These devices are connected to an edge gateway and publish runtime process and status data to a cloud-based data historian.

Building upon IP infrastructure enables global access to HART data. However, using HART-IP to pierce the barriers to enterprise integration must be done carefully with well-implemented security practices. Building upon TLS/DTLS provides proven communication security for a wide range of applications. HART-IP extends this by allowing specific users to have read-only access. This allows (for example) data historians to be effectively connected via a "data Diode".

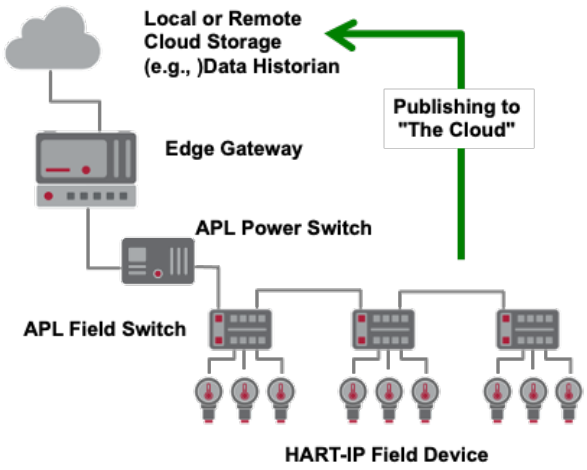

Figure 6. HART-IP Monitoring Network

HART-IP specifies minimum standardized security. Security for IP-based products continues to evolve and HART-IP is designed to benefit from these security measures. Consequently, field upgrade support is mandatory for all HART-IP devices. Additional security protocols and measures (like VPN, firewalls, intrusion detection systems, etc.) can be added to these minimums as required by the users' application or as options to the manufacturers' offering. A system can thus be designed to provide multiple layers of encryption, authentication, authorization, and verification to provide defense-in-depth security.

HART-IP provides excellent tools to ensure cyber security. This is a very good start, but users must recognize that plant policies and procedures are just as important to the plants' operational security, as are the minimum security features available in all HART-IP devices.

ANNEX A. GLOSSARY

CEF	**C**ommon **E**vent **F**ormat (used in syslogging).
Client	The requesting program or user in a client/server relationship. Initiates requests on behalf of a user/program and waits for replies to transfer the requested information back to the user/program.
DTLS	**D**atagram **T**ransport **L**ayer **S**ecurity
Encryption	Encryption is a two-way function where information is encrypted allowing the information to be decrypted later. The encrypted data is scrambled and unreadable by third parties.
IANA	**I**nternet **A**ssigned **N**umbers **A**uthority
Interoperability	Interoperability is the ability for like devices from different manufacturers to work together in a system and be substituted one for another without loss of functionality at the host system level.
IP	**I**nternet **P**rotocol – a data-oriented, unreliable network layer protocol used for communicating data across a packet-switched network. (Network Layer of the OSI model.
IP Address	The internet protocol address of the device.
LDH	**L**etters (a through z; A through Z), **D**igits (0 through 9), and **H**yphens. The only characters allowed in an Internet HOSTNAME.
Network Device	A device with a direct Physical Layer connection to the network. Each network device (e.g., field device or gateway) has a HART Unique Address that is used in communication with the device.
Node	An addressable logical or physical device attached to the network.
NTP	**N**etwork **T**ime **P**rotocol
Packet	A generic reference to the set of data communicated across a network.
PAM	**P**lant **A**sset **M**anagement
PDU	**P**rotocol **D**ata **U**nit. The packet of information being communicated.
Physical Layer	Physical Layer is responsible for transmission of the raw bit stream and defines the mechanical and electrical connections and signaling parameters for devices.
Port	The specific UDP or TCP port used for connection of this service.
Security Manager	An application that manages the Network Device's security resources and monitors the status of the network security.
Server	Provides services to other programs or users. Never initiates requests but may issue unsolicited responses (publish process data or notifications). The server waits for requests from one or more clients and replies to client with appropriate response.
Session	A semi-permanent interactive information exchange.
SIEM	**S**ecurity **I**nformation and **E**vent **M**anagement. A syslog server with analysis capabilities.
Socket	One endpoint of a two-way communication link between two programs running on the network. The communication link uses either UDP or TCP as the message protocol, a local and remote IP address and a local and remote port.
Syslog Server	A syslog Server is a central destination and repository on a network that receives syslog messaged emitted by syslog clients (e.g., HART-IP field devices). Most syslog Servers include a database that allows the syslog messages to be reviewed and anomalous events traced.
TCP	**T**ransmission **C**ontrol **P**rotocol – reliable, connection-oriented, in-order delivery of a stream of bytes.
TLS	**T**ransport **L**ayer **S**ecurity
UDP	**U**ser **D**atagram **P**rotocol – does not guarantee reliability or packet ordering and is connectionless.

ANNEX B. REFERENCES

[1] FieldComm Group, https://fieldcommgroup.org

[2] HART Communications Protocol Specifications. https://www.fieldcommgroup.org/hart-specifications

[3] The Transport Layer Security (TLS) Protocol Version 1.2, https://tools.ietf.org/html/rfc5246

[4] The Transport Layer Security (TLS) Protocol Version 1.3, https://tools.ietf.org/html/rfc8446

[5] Guidance for External PSK Usage in TLS, https://www.rfc-editor.org/rfc/rfc9257

[6] Stouffer; Lightman; Pillitteri; Pease; Tang; Zimmerman. "Guide to Operational Technology (OT) Security". NIST. https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-82r3.ipd.pdf

