



SANS Institute

Information Security Reading Room

HL7 Data Interfaces in Medical Environments: Understanding the Fundamental Flaw in Healthcare

Dallas Haselhorst

Copyright SANS Institute 2021. Author Retains Full Rights.

This paper is from the SANS Institute Reading Room site. Reposting is not permitted without express written permission.

HL7 Data Interfaces in Medical Environments: Understanding the Fundamental Flaw in Healthcare

GIAC (GSEC) Gold Certification

Author: Dallas Haselhorst
Email: dallas@treetopsecurity.com / Twitter: @oneoffdallas
Advisor: Sally Vandeven
Accepted: August 2017

Abstract

Ask healthcare IT professionals where the sensitive data resides and most will inevitably direct attention to a hardened server or database with large amounts of protected health information (PHI). The respondent might even know details about data storage, backup plans, etc. Asked the same question, a penetration tester or security expert may provide a similar answer before discussing database or operating system vulnerabilities. Fortunately, there is likely nothing wrong with the data at that point in its lifetime. It potentially sits on a fully encrypted disk protected by usernames, passwords, and it might have audit-level tracking enabled. The server may also have some level of segmentation from non-critical servers or access restrictions based on source IP addresses. But how did those bits and bytes of healthcare data get to that hardened server? Typically, in a way no one would ever expect... 100% unencrypted and unverified. HL7 is the fundamentally flawed, insecure standard used throughout healthcare for nearly all system-to-system communications. This research examines the HL7 standard, potential attacks on the standard, and why medical records require better protection than current efforts provide.

1. Introduction

In comparison to other industries, healthcare contains some of the most sensitive and valuable data to an attacker. Although the cost for a stolen electronic health record (EHR) on the dark web has dropped from previous years, complete EHR databases have recently sold for as much as \$500,000 (Chickowski, 2017). EHRs are highly valuable to criminals because individual records and the databases that contain them are a wealth of information. According to Fuentes, “EHR data is unique in a way that it includes protected identifiable information (PII), along with medical, insurance, and financial information” (2017). Figure 1 depicts the sale of individual medical records on a dark web marketplace.

The screenshot shows a marketplace listing for 'Medical Fullz'. The listing includes a cartoon illustration of a doctor, a list of data fields (e.g., Patient ID, Name, Address, Insurance), and a table of features. The purchase price is listed as USD 5.00, and there are 'Buy Now' and 'Queue' buttons. The listing also shows a quantity of 1 and a default shipping time of 1 day.

Features	Features
Product class	Digital goods
Quantity left	Unlimited
Ends in	Never
Origin country	Worldwide
Ships to	Worldwide
Payment	Escrow

Figure 1: Medical records for sale on the now defunct AlphaBay (dark web), image source - Deep.Dot.Web

Healthcare essentially becomes an all-in-one stop for data. With all the data in one location, this opens the door to a wide range of illicit opportunities on the black market and dark web. Medical records contain full names, physical addresses, social security numbers, driver’s license information, email addresses, and dates of birth. One can even frequently find answers to common secondary security questions, e.g., past addresses and

Dallas Haselhorst

phone numbers, place of birth, relatives (and related information), employer information, and much more. If a cybercriminal is interested in running scams for fraudulent tax returns, healthcare has it. If an attacker wants to steal an identity and sign-up for credit cards, healthcare has information for that too. Healthcare even has all the data necessary to perform fake prescription refills. A single EHR contains the data for any of these illegal activities.

Healthcare is known for its regulation and significant penalties for data breaches. Given the rising costs associated with a breach and the value of stolen healthcare data, one would assume it is extremely well-protected. Unfortunately, that could not be further from the truth. Healthcare is considered abysmally behind in its cybersecurity efforts: “The health care industry definitely lags all other industries, except maybe higher education” (Bur, 2016). Depending on the measurement methodology, some estimates place healthcare firmly behind other critical industries by as much as five to ten years.

Criminal organizations are targeting healthcare data more than ever and there is a somewhat lackadaisical approach to securing medical data. Meanwhile, there is a strong desire to connect everything related to healthcare causing additional burdens. Connectivity has proved to be a double-edged sword healthcare has only started to understand. As stated by the Health Care Industry Cybersecurity Task Force, “Patients and physicians have derived many benefits from EHRs including giving patients the ability to access their information through portals and giving providers the ability to more easily share patient information. However, this digitization resulted in an increased attack surface...” (2017). Healthcare has arguably felt the need for data sharing and data connectivity far more and far longer than any other industry. The prolonged and ever-increasing connectivity has added many legacy standards, applications, and devices over the years. The past and present push for more connectivity could add layers of difficulty to securing healthcare environments because it is a constantly moving target.

2. Role of HL7 in Healthcare

Understanding the importance of safely handling medical data and providing secure connectivity in healthcare is not complete without a discussion of the HL7

standard. However, in most discussions, HL7 is entirely forgotten because it operates behind the scenes. Health Level 7 (HL7) provides "a framework (and related standards) for the exchange, integration, sharing, and retrieval of electronic health information" (Introduction to HL7 Standards, 2017). HL7 is the underlying standard that provides a common syntax for systems to interoperate with one another and share information. It is important to note that there are multiple versions of HL7, namely version 2 and version 3, with a newer standard on the horizon. Version 3 does offer many improvements, although it is not backwardly compatible. For this reason and many more, version 3 never supplanted version 2 and it is still not as widely adopted (Corepoint Health, 2017). Consequently, this research focuses on version 2 of the HL7 standard although similar security issues exist in version 3 as well.

The goal of healthcare IT has long been the exchange of patient health data and the ability to send and share data accurately with other disparate systems. This data transfer is important if a patient's insurance or address has changed, but also if a bed number changes during a hospital stay. It ensures patient data is accurate throughout a visit while also providing speed in patient care. An example of speed in patient care is the drastically reduced time to transfer radiology and lab results from ancillary systems to the primary system most hospital personnel use. HL7 communications assist in this function by eliminating most manual data entry, facilitating the delivery of time sensitive information and reducing the number of errors.

It is rare, if not impossible, to find a single hospital system that performs all the necessary medical functions from beginning to end. Even a small, limited care facility will typically have several systems to support patient care or back-end business tasks. For example, a small facility might have an ancillary system for pharmacy functionality in addition to their primary system which performs a bulk of the hospital duties. In a way, HL7 provides the two systems with a similar language to communicate with one another. This standard means of communications facilitates the data flow between two systems as seen in Figure 2, thereby allowing the overall hospital to function more efficiently.




Figure 2: Simple point-to-point HL7 interface

If a small facility does not have an in-house lab, it may have an HL7 feed to an external lab to eliminate manual data entry or the faxing of patient data and orders. Even in a small environment, it simply is not possible for someone to manually input data into each of those systems. In a slightly larger hospital, the facility will likely have a financial/admissions system, a few drug/supply dispensing machines in patient areas, a lab system, and a separate radiology system. Rather than create simple point-to-point interfaces, the hospital may also incorporate an interface engine to assist with the data interchange between these systems. An interface engine (aka integration engine) is a go-between for different systems that monitors different types of interfaces and communication points and performs actions according to rules defined by the organization (Orion Health, 2017). As part of its core functionality, an interface engine replicates HL7 messages to multiple systems while formatting the messages to meet the needs of individual receiving systems.

Suddenly, the simple two-system configuration described previously has grown significantly. This growth continues as the relative size of a hospital and the number of services increases. As represented in Figure 3, a large facility could require the primary admitting system to send patient record data to

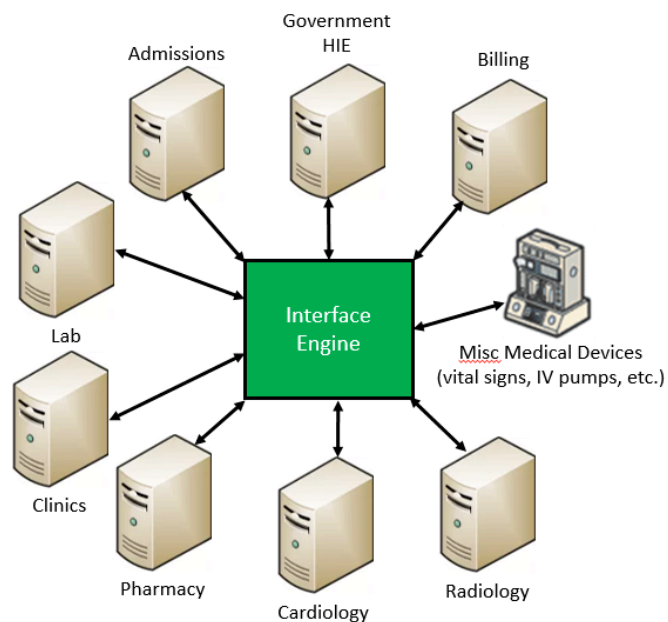


Figure 3: Larger installation with an HL7 interface engine

several systems – lab, billing, radiology, pharmacy, and maybe even a government health information exchange (HIE). HL7 data may also flow back from any of those systems or miscellaneous medical devices (such as those used for collecting vital signs, IV pumps, or other patient centric data) to the clinician charting application.

3. HL7 Messages

An HL7 interface connects systems, however, HL7 data is not typically flowing between them non-stop. Instead, HL7 messages are sent on an as-needed basis when an event triggers a message. The contents of HL7 messages vary wildly even though the standard establishes message structure to some degree. Every HL7 message begins with a message header (MSH) segment and it "defines the message's source, purpose, destination, and certain syntax specifics" (HL7 MSH – Message Header, 2017). An HL7 message itself is little more than text with field delimiters as shown in Figure 4 below.

```
MSH|^~\&|SENDING_APPLICATION|SENDING_FACILITY|RECEIVING_APPLICATION|RECEIVING_F
ACILITY|20110613083617||ADT^A01|934576120110613083617|P|2.3||||
```

Figure 4: Example message header (MSH)

In the example MSH segment, one can see the delimiters in use immediately following the 'MSH' text. The message used the standard and recommended values – '|' is the field separator, the '^' is the component separator, the '~' is the repetition separator, the '\' is the escape character, and the '&' is the sub-component separator. Following the delimiter field are the sending and receiving fields, a date/time field, the type of event, and other miscellaneous information. The "required" data fields found both inside and outside of the MSH segment are heavily dependent on the receiving system and the HL7 standard itself.

3.1. ADT Messages

ADT (admissions, discharges, and transfers) messages are one of the most widely used HL7 message types. These messages are primarily for patient administration. ADT messages carry patient demographic information as well as information about patient visits/encounters. When a hospital admits a new patient, an outbound ADT message will

most likely trigger although this is once again dependent on the needs of receiving systems. More specifically, if the patient visit is a standard admission, then an ADT-A01 message much like the example in Figure 5 is sent. The A01 is an event type within the ADT message type that specifies it is an admission and not a pre-admission, discharge, etc. If the patient switches to a new room or bed, an A02 (transfer) or A08 (patient update) message might get triggered instead. There are numerous other ADT message types and they can cover a variety of activities including patient registrations, changing an inpatient to an outpatient, or canceling events.

```
MSH|^~\&|SENDING_APPLICATION|SENDING_FACILITY|RECEIVING_APPLICATION|RECEIVING_F
ACILITY|20110613083617||ADT^A01|934576120110613083617|P|2.3|||
EVN|A01|20110613083617||
PID|1||135769||MOUSE^MICKEY^||19281118|M|||123 Main St.^Lake Buena Vista^FL
^32830|(407)939-5555^^^ohtoodles@notdisney.com||||1719|999999999||
|MOUSETOWN|||||
NK1|1|MOUSE^MINNIE|WIFE||||NK
PV1|1|O||||^^^^^^^^^^|^^^^^^^^^^
AL1|1|^Penicillin||Anaphylactic shock
AL1|2|^Cat dander||Skin rash
```

Figure 5: Example ADT message

Different HL7 message types have structural commonalities and contain similar, sensitive data. In the ADT example, the message has an MSH segment, but it also has an EVN (event) segment and a PID (patient identifier) segment. Many of the HL7 message types will have the EVN and PID segments as well as other required segments. The HL7 standard guides the data structure for the various message types. Table 1 below shows some "security significant" ADT data fields along with their segment and field number followed by the information contained in those fields from the ADT example.

Segment	Field #	Field Definition	Field Contents
PID	5	Patient Name	MOUSE^MICKEY
PID	6	Mother's Maiden Name	<i>NOT SPECIFIED</i>
PID	7	Date/Time of Birth	19281118

PID	11	Patient Address	123 Main St.^Lake Buena Vista^FL^32830
PID	13.1	Phone Number – Home	(407)939-5555
PID	13.4	Email Address	ohtoodles@notdisney.com
PID	19	SSN Number	999999999
PID	20	Driver’s License Number	<i>NOT SPECIFIED</i>
PID	23	Birth Place	MOUSETOWN
NK1	2	Next of Kin Name	MOUSE^MINNIE
AL1 (1)	3	Allergy	Penicillin
AL1 (1)	4	Allergic Reaction	Anaphylactic shock
AL1 (2)	3	Allergy	Cat dander
AL1 (2)	4	Allergic Reaction	Skin rash

Table 1: Sensitive data found in a standard HL7 message

3.2. ORM Messages

ORM (order) messages are general orders that can refer to materials or services, for example, ordering a drug such as morphine or requesting a blood screen. An ORM message, as previously indicated, shares many commonalities with other HL7 messages. However, a segment only found in order-related messages is the ORC segment. The ORC segment specifies order details and can include fields such as who ordered the test/drug, who entered the information, quantities, etc. ORM messages may also have OBR segments detailing additional order information. Figure 6 is an example ORM message. For clarification, some pharmacy dispensing systems require a pharmacy/treatment encoded order (RDE) message for drug ordering versus a standard ORM message.

```
MSH|^~\&|SendingApp|SendingFac|ReceivingApp|ReceivingFac|20120411070545||ORM^00
1|59689|P|2.3
PID|1|12345|12345^^^MIE&1.2.840.114398.1.100&ISO^MR||MOUSE^MICKEY^S||19281118|M
|||123 Main St.^Lake Buena Vista^FL^3283||||||||||||||||
```

Dallas Haselhorst


```

OBX|4|NM|monos^Monos^Local^5905-5^Monos^LN||16||%|4-
15|H|||F|||20120410160227|lab|12^XYZ LAB|
OBX|5|NM|eo^Eos^Local^713-8^Eos^LN||3||%|0-8|||F|||20120410160227|lab|12^XYZ
LAB|
OBX|6|NM|baso^Baso^Local^706-2^Baso^LN||0||%|0-
1|||F|||20120410160227|lab|12^XYZ LAB|
OBX|7|NM|ig^Imm Gran^Local^38518-7^Imm Gran^LN||0||%|0-
2|||F|||20120410160227|lab|12^XYZ LAB|
OBX|8|NM|rbc^Rbc^Local^789-8^Rbc^LN||4.02|/pl|4.07-
4.92|L|||F|||20120410160227|lab|12^XYZ LAB|
OBX|9|NM|hgb^Hgb^Local^718-7^Hgb^LN||13.7|g/dl|12.0-
14.1|||F|||20120410160227|lab|12^XYZ LAB|
OBX|10|NM|hct^Hct^Local^4544-3^Hct^LN||40||%|34-
43|||F|||20120410160227|lab|12^XYZ LAB|
OBX|11|NM|mcv^Mcv^Local^787-2^Mcv^LN||80|fl|77-
98|||F|||20120410160227|lab|12^XYZ LAB|
OBX|12|NM|mch^Mch||30|pg|27-35|||F|||20120410160227|lab|12^XYZ LAB|
OBX|13|NM|mchc^Mchc||32|g/dl|32-35|||F|||20120410160227|lab|12^XYZ LAB|
OBX|14|NM|plt^Platelets||221|/nl|140-400|||F|||20120410160227|lab|12^XYZ LAB|

```

Figure 7: Example ORU message

3.4. ACK Messages

While there are several other message types including well-known ones such as SUI (for scheduling) and DFT (detailed financial transactions for billing), the last type discussed here is ACK (acknowledgment) messages. Although HL7 communication does not always use ACK messages, understanding ACK usage is essential because it can cause messaging failures if excluded. In some cases, an ACK is sent immediately after a message is received. There are instances where return messages are only sent *after* the receiving system has 1) succeeded in processing the message (ACK), 2) received the message, but processing failed such as the error message below (also an ACK, but with description), 3) or failed processing (NACK). Another outcome is the lack of any returned messages. Failure to receive any indications back can mean the receiving system is down or not processing messages at that time.

Success

```

MSH|^~\&|SENDING_APPLICATION|SENDING_FACILITY|RECEIVING_APPLICATION|RECEIVING_F
ACILITY|20110614075841||ACK|1407511|P|2.3|||

```

```
MSA|AA|1407511|Success||
```

Error

```
MSH|^~\&|SENDING_APPLICATION|SENDING_FACILITY|RECEIVING_APPLICATION|RECEIVING_FACILITY|20110614075841||ACK|1407511|P|2.3|||||
MSA|AE|1407511|Error processing record!||
```

Figure 8: Example ACK and failure messages

For a more complete listing of HL7 message types and field-by-field specifications, please reference <https://corepointhealth.com/resource-center/hl7-resources/hl7-messages>.

4. Insecurity Built-In

The overview on the value of stolen health records and the data contained in HL7 messages helps convey the need for security in healthcare. Regrettably, HL7 could not be any further from addressing that need. The HL7 standards body states, “In the Security TC we have assumed that encryption happens below the application layer, e.g., via IPsec or TLS, not within HL7 messages” (HL7 International, 2007). HL7 International created the standard with a strict focus on the OSI (Open Systems Interconnection) application layer, leaving security as an implementation detail. This concept means that despite the sensitivity of the data, HL7 is sent and received as clear-text with the *potential* of encryption between the endpoints.

The lack of encryption in HL7 and the relative insecurity of healthcare coincides with recent survey results from HIMSS (Health Information and Management Systems Society). As depicted in Figure 9, the *HIMSS Cybersecurity Report* estimates only 64.0% of U.S. based provider organizations (e.g., hospitals, physician offices) encrypted data in transit (n.d.). It is highly possible this number may have been increased by survey respondents either overlooking the usage of HL7 or the belief that a VPN (virtual private network) or some other means is securing the data.

	Total	Acute Care	Non-Acute	DIFF
Antivirus/malware	86.0%	84.9%	90.3%	-5.4%
Firewalls	80.7%	78.2%	90.3%	12.2%
Data encryption (data in transit) ←	64.0%	68.1%	48.4%	19.7%
Audit logs of each access to pt. health and financial records	60.0%	59.7%	61.3%	-1.6%

Figure 9: HIMSS Cybersecurity Report - Data Encryption

Overall, the transmission of HL7 messages is similar to the insecurities found in FTP (file transfer protocol) and older implementations of SMTP (simple mail transfer protocol). Somewhat ironically, even today interface engines often utilize FTP to perform batch transfers of HL7 messages. Individual HL7 messages or batch transfers are also still sent using SMTP. A conventional store and forward interface example is the accumulation of individual financial (DFT) messages. Following the message collection, the single file gets sent to a billing system (possibly external to the network) via a daily FTP batch. The *HIMSS Cybersecurity Report* did not detail why the data in transit numbers were only 64.0%. One possibility is that while HL7 might fly under the “not encrypted” radar, FTP is a promising reason for why survey respondents chose “no” when asked if their data was encrypted. Standard FTP has well-known insecurities and if it exists in an environment, the IT department staff is most likely aware of it. Without further clarification of how the question was asked or the technical knowledge level of the person answering the questions, it is impossible to know if the survey results are high or low. It also raises the issue of whether HL7 or FTP were taken into consideration when answering the questions.

It might come as some surprise that HL7 provides an even lower level of security than FTP. Unlike FTP, HL7 does not require any authentication nor is there any way to include authentication when establishing a connection. FTP is avoided entirely in most environments due to numerous security concerns. Meanwhile, an overwhelming majority of healthcare facilities use HL7 for some of the most sensitive data imaginable even though it is a security nightmare. As shown in the two configuration screenshots below, a standard HL7 interface is nothing more than a destination IP address and TCP port on a source system with a corresponding listening port on a destination system.

Connector Type: **TCP Sender** Wait for previous destination

Destination Settings

Queue Messages: Never On Failure Always

Advanced Queue Settings: Retries

Validate Response: Yes No

Reattach Attachments: Yes No

TCP Sender Settings

Transmission Mode: **MLLP**

MLLP Sample Frame: <VT> <Message Data> <FS> <CR>

Remote Address: **10.0.0.127**

Remote Port: **6661**

Figure 10: A TCP sender configuration (IP address and port) in Mirth, a widely used interface engine

Edit Channel -

Summary | Source | Destinations | Scripts

Connector Type: **TCP Listener**

Listener Settings

Local Address: All interfaces Specific interface:

Local Port: **6661**

Figure 11: A TCP listener port in Mirth

5. Man-In-The-Middle (MITM) Attacks

As with any insecure messaging, there are numerous opportunities to attack the underlying communications. If the attacker can gain access to the same network, the possibility for ARP (address resolution protocol) poisoning or ARP spoofing is extremely high. ARP spoofing is relatively easy to perform and it allows an attacker to intercept all communications between two systems using a packet sniffer. This process is also known as eavesdropping since neither the sending nor receiving system is aware of the intermediate system or redirection. In Figure 12, an attacker uses ARP spoofing to intercept the communications between an interface engine and a lab system. This MITM technique allows an attacker to gather PHI or even modify network traffic in real-time with potentially fatal outcomes.

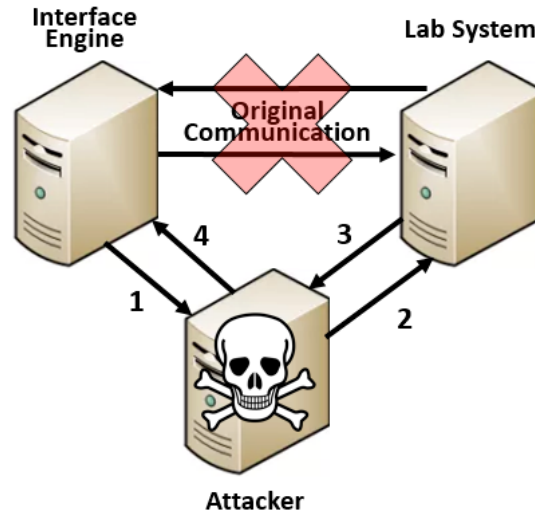


Figure 12: ARP Spoofing Attack of HL7

An attacker may establish a foothold in an environment to quietly collect stolen medical record data and eventually use it for identity fraud or other financial gains. As shown in Figure 13, when used in conjunction with a MITM attack the captured HL7 message is easily recognizable when displayed in Wireshark. After the data capture phase, the data could be sent off-site periodically and the collection process restarted. If no service disruptions occurred with the actual HL7 interfaces, this data exfiltration could conceivably continue for weeks, months, or even years without detection. The relative size of the facility and the number of patient encounters are the only limiting factors in the number of collected records.

```

0000 00 0c 29 7f 16 f3 00 50 56 2f 29 76 08 00 45 00 ..)....P V/)v..E.
0010 02 02 4e 58 40 00 80 06 95 a1 0a 00 00 7e 0a 00 ..NX@... ..~...
0020 00 7f e6 6c 1a 05 e5 3c af 9a d0 5a a5 12 50 18 ...l...< ...Z..P.
0030 01 00 a3 64 00 00 0b 4d 53 48 7c 5e 7e 5c 26 7c ...d...M SH|^~\&|
0040 53 45 4e 44 49 4e 47 5f 41 50 50 4c 49 43 41 54 SENDING_ APPLICAT
0050 49 4f 4e 7c 53 45 4e 44 49 4e 47 5f 46 41 43 49 ION|SEND ING_FACI
0060 4c 49 54 59 7c 52 45 43 45 49 56 49 4e 47 5f 41 LITY|REC EIVING_A
0070 50 50 4c 49 43 41 54 49 4f 4e 7c 52 45 43 45 49 PPLICATI ON|RECEI
0080 56 49 4e 47 5f 46 41 43 49 4c 49 54 59 7c 32 30 VING_FAC ILITY|20
0090 31 31 30 36 31 33 30 38 33 36 31 37 7c 7c 41 44 11061308 3617||AD
00a0 54 5e 41 30 31 7c 39 33 34 35 37 36 31 32 30 31 T^A01|93 45761201
00b0 31 30 36 31 33 30 38 33 36 31 37 7c 50 7c 32 2e 10613083 617|P|2.
00c0 33 7c 7c 7c 7c 0d 45 56 4e 7c 41 30 31 7c 32 30 3|||.EV N|A01|20
00d0 31 31 30 36 31 33 30 38 33 36 31 37 7c 7c 7c 0d 11061308 3617|||.
00e0 50 49 44 7c 31 7c 7c 31 33 35 37 36 39 7c 7c 4d PID|1||1 35769||M
00f0 4f 55 53 45 5e 4d 49 43 4b 45 59 5e 7c 7c 31 39 OUSE^MIC KEY^||19
0100 32 38 31 31 31 38 7c 4d 7c 7c 7c 31 32 33 20 4d 281118|M |||123 M
    
```

Figure 13: Unencrypted HL7 data gathered from Wireshark on an attacking MITM system

The motivations for other attackers may go beyond simply collecting data. Another potential attack involves the manipulation of HL7 data. A cybercriminal could remove entire HL7 DFT messages to eliminate billing charges. The attacker could also cause mass disruption by swapping patient names on patient-viewed results or billing statements. Technically speaking, each of those “minor” mix-ups is a data breach requiring substantial legwork to correct. Other consequences may include fines from the breach and the loss of patient confidence in the facility.

A malicious attacker could also make changes with far more nefarious intentions causing bodily harm or even death. An attacker could modify simple HL7 data such as the weight of a patient. An incorrect weight sounds harmless enough during a routine check-up, but it could drastically change the dosing for numerous drugs associated with ongoing treatments. Another attack might result in the improper dosing of a chemotherapy drug or additional drug orders before and after a shift change. It is unlikely the receiving system would alert technicians or pharmacists when requested dosages are over a certain limit or prescribed more frequently than usual. Falsified test results could also drastically change a patient’s care plan. Very rarely, if ever, would anyone identify this type of data tampering and thus, it is currently the responsibility of hospital staff to manually detect these anomalies.

Consider the hypothetical scenario of a high-level politician targeted by a nation-state or high-level cybercriminal organization. It is not outside the realm of possibility for one of those bad actors to attempt an assassination wielding HL7 as a weapon. As described earlier, HL7 ADT messages have allergy segments and HL7 also sends drug orders. Now consider the target was severely allergic to penicillin leading to anaphylactic shock. When admitted to the hospital, the attacker could modify the ADT AL1 segment to read “no known allergies” for the target. The attacker could follow-up the falsified ADT message with an order for penicillin. This attack scenario goes beyond a single penicillin allergy as there are thousands of other allergic reactions or drug interaction possibilities to consider as well.

6. HL7 On Life Support

The de facto standard healthcare has used for years to communicate is not going to change overnight. It is not as if there has been a competing standard for the past 30 years where only a handful of interfaces are affected. Even if the industry moved to a new and secure standard tomorrow, there are hundreds of thousands of HL7 interfaces in operation today. Security-oriented changes to HL7 and other insecure transports would allow their continued usage without a complete replacement. Improving the security of existing interfaces could save a significant amount of cost and time while still protecting sensitive data.

First, encryption must be deployed wherever possible, especially in instances where health data is traversing different network segments or the Internet. Most often, an encrypted solution will involve sending the data over a VPN tunnel. In the case of an actual point-to-point VPN tunnel, this may be all that is needed. An SSH tunnel could provide similar security gains as a VPN. Other benefits of an SSH tunnel include less vendor involvement, fewer configuration changes, and an easier setup.

If a VPN or an SSH tunnel is not a viable option, another possibility is the use of network segmentation to help secure HL7. While network segmentation will almost always improve security, it does not necessarily eliminate the risk. Proper segmentation is also notoriously difficult for interface engines. By their very nature, interface engines must communicate with numerous outlying systems. This communication requirement often results in interface engines sitting on a “lowest common denominator” network segment, i.e., a network segment that can communicate with nearly everything on an internal hospital network. The problem with this approach is a simple email phishing attack or browser-based malware infection can give an attacker access to valuable medical record data. Data handled by the interface engine is far too sensitive for this to occur and at the very least, the interface engine must reside on a different network segment than standard user workstations.

Applying static ARP entries is another, under-utilized means of preventing HL7 MITM attacks. As with other HL7 fixes, static ARP entries will likely involve vendors, change control, and will have budgetary impacts. A static ARP entry would ensure the

Dallas Haselhorst

systems could only communicate with media access control (MAC) address tied to a related IP address. This defense mechanism would prevent a rogue system on the same network from gaining a foothold via ARP spoofing. However, static ARP entries would do little to protect the sensitive data in transit as the interfaces would still communicate as clear-text. Thus, if an attacker gained access to network hardware such as a router or switch between the HL7 interface sender and receiver, the data would pass unencrypted and completely visible.

Re-architecting any working solution in the name of improving security is often met with backlash and interfaces are no exception. Some complex interfaces were multi-month or multi-year efforts. In many cases, the interface engineer who developed the interfaces may not be with the organization anymore. Even simple interface modifications such as changing an IP address can involve a substantial amount of planning. The potential for downtime also exists if someone forgets a minor technical detail in the process. Despite the roadblocks and possible missteps, securely re-architecting existing HL7 solutions must be performed to secure the data and limit would-be attackers access to it.

7. What's Next?

Healthcare needs a better means of system-to-system communication. After all, the most frequently used version of HL7, version 2, was created in 1989 (HL7 Frequently Asked Questions, 2017). Fortunately, there is a new standard known as Fast Healthcare Interoperability Resources (FHIR) and it promises to revolutionize how healthcare data is shared. The creators of the HL7 standard, Health Level Seven International, also created the FHIR standard. FHIR uses an application programming interface (API) that allows for much easier data access. It achieves this via an HTTP-based RESTful protocol in either XML or JSON formats.

FHIR has its benefits, but corrections to underlying issues must occur before its widespread adoption. In the FHIR security section, the HL7 organization states “TLS/SSL *SHOULD* be used for all production data” (HL7.org, 2017). The problem is *should* is not synonymous with *required* or *must*. The standard further distances itself

from guaranteeing encrypted traffic by also stating “the service base URL will specify *whether* SSL is required” (HL7.org, 2017). Securely implementing FHIR does not have the appropriate guidance from the standards body. When added to the increased complexity of configuring SSL/TLS certificates and the reduced ability to troubleshoot network communications, healthcare may see security eliminated from real-world FHIR implementations before it gets off the ground. Undoubtedly, the lack of built-in security measures will lead to what the IT and InfoSec industries have seen time and time again, i.e., the easy install of today will become the gaping security hole of tomorrow.

Healthcare itself and various technical standards bodies have not adequately answered this challenge up to this point. Thus far, government intervention has fallen short as well. Most healthcare professionals recognize HIPAA (Health Insurance Portability and Accountability Act) as the attempt of the government to steer the healthcare ship in the United States. Right or wrong, much of HIPAA has previously focused on patient privacy and policy instead of deep-seated technical security issues. That stance changed to some degree in the Omnibus Security Rule update in 2013.

Still, the guidance found in the revised rule was left open to interpretation and those interpretations can vary wildly. For instance, something as simple as workstation security has no implementation specifics and instead says workstations should be protected from unauthorized use (AHIMA Practice Brief, 2013). The rule does not specify if a username and password are sufficient or if secondary measures such as biometrics or two-factor authentication (2FA) are needed. After all, it is reasonable to believe a malicious insider or an outside attacker could gain *unauthorized* access via stolen user credentials. The rule also does not specifically mention the use of full-disk encryption (FDE) for data at rest, which would further protect against unauthorized use. While it is common to use FDE on laptops to protect data if they are lost or stolen, FDE on servers is infrequently used or recommended in healthcare. The lack of understanding where sensitive data is stored or the lack of regulatory requirements may underscore this oversight. The important takeaway is that despite interface engines receiving and storing massive amounts of PHI, they do not draw the same level of scrutiny as other devices where PHI is not stored.

Dallas Haselhorst

Whether discussing data in transit or at rest, the HIPAA Security Rule falls short by simply stating that encryption be addressable rather than required (Snell, 2015). Essentially, if an organization cannot encrypt the data, someone should be prepared to explain why it is not feasible. The rule has shifted more toward required over time, yet required is still not mandatory in all instances. With regards to data in transit, the HL7 standard does not support encryption. The FHIR standard does not require SSL/TLS so a vendor may not support it, thereby preventing the organization from utilizing it. If a breach occurred, it is unclear whether those prior justifications would satisfy the “addressable, not required” guidance from the HIPAA Security Rule.

8. Conclusion

Healthcare communications such as HL7 need security measures added for both existing and new installations. Its successor, FHIR, should require encryption to achieve compliance with the standard. The bottom line is that for sensitive medical data, the question of whether the communications are secure or not should not vary from one implementation to another. Protecting medical record data or any medical device is nothing more than a pipedream without secure communications as the standard in healthcare. As it is today, HL7 will continue to be the “data privacy hell of communication” as security researcher, Hannes Molsen, described it to the FDA and other U.S. government agencies (Capital Reporting Company, 2016). Furthermore, discussion on not only device security, but the security of inter-system communications is necessary to move the industry forward and ensure “we will have a safe system, not only safe boxes in an insecure system” (Capital Reporting Company, 2016).

In many ways, outside forces including the government, patients longing for access to more data, and even healthcare itself have unintentionally created the perfect storm placing everyone’s data at risk. The perils associated with not securing healthcare data are endless and that will not change in our data-driven, connected world. The value of medical records is simply too high for criminals to cease their assault on healthcare and the potential usage of HL7 and EHRs for malicious purposes is far too great to ignore. At a minimum, the healthcare industry needs to focus on improving security across the board and achieving the same level of security as other critical industries –

Dallas Haselhorst

finances, utilities, and retail. Without significant improvements, healthcare will continue to be nothing more than low-hanging fruit with high rewards.

As with any issue and particularly those in information security, this challenge will only get addressed once light shines on the problem. Healthcare needed better security yesterday and change needs to happen. The security woes found in healthcare were not created overnight and there are no quick and easy solutions either. Healthcare communications standards must have security built-in so the new implementations of tomorrow, next week, and next year can realize significant improvements day one rather than as an afterthought. The healthcare industry will require time to make long-lasting improvements and those advances need to start now. Medical data is far too valuable to continue neglecting security.

References

- AHIMA Practice Brief. (November 2013). *HIPAA Security Rule Overview (2013 update)*. Retrieved from <http://library.ahima.org/doc?oid=300262>
- Bur, Jessie. (2016, August 23). Health Care Cybersecurity A Decade Behind, HIMSS Survey Reveals. Retrieved April 4, 2017, from <https://www.meritalk.com/articles/health-care-cybersecurity-a-decade-behind-himss-survey-reveals/>
- Capital Reporting Company. *Moving Forward: Collaborative Approaches to Medical Device Cybersecurity*. (2016, January 20). Retrieved from <https://www.fda.gov/downloads/MedicalDevices/NewsEvents/WorkshopsConferences/UCM489249.pdf>
- Chickowski, Ericka. (2017, February 21). Stolen Health Record Databases Sell For \$500,000 In The Deep Web. Retrieved June 2, 2017, from [http://www.darkreading.com/attacks-breaches/stolen-health-record-databases-sell-for-\\$500000-in-the-deep-web/d/d-id/1328225](http://www.darkreading.com/attacks-breaches/stolen-health-record-databases-sell-for-$500000-in-the-deep-web/d/d-id/1328225)
- Corepoint Health. Versions of the HL7 Standard. Retrieved July 2, 2017 from <https://corepointhealth.com/resource-center/hl7-resources/hl7-standard-versions>
- Fuentes, Mayra Rosario. (2017). *Cybercrime and other Threats faced by the Healthcare Industry*. Retrieved from <https://www.trendmicro.com/content/dam/trendmicro/global/en/security-intelligence/research/reports/wp-cybercrime-&-other-threats-faced-by-the-healthcare-industry.pdf>
- Health Care Industry Cybersecurity Task Force. (June 2017). *Report on Improving Cybersecurity in the Health Care Industry*. Retrieved from <https://www.phe.gov/Preparedness/planning/CyberTF/Documents/report2017.pdf>
- HIMSS. *2016 HIMSS Cybersecurity Study*. (n.d.). Retrieved from http://www.himss.org/sites/himssorg/files/images/HIMSS_CybersecurityReport_2016.pdf
- HL7 Frequently Asked Questions. Retrieved June 1, 2017 from <http://www.hl7.org/about/FAQs/index.cfm?ref=nav>

HL7 International. (2007, August 31). Implementation FAQ: Encryption and Security.

Retrieved January 26, 2017, from

http://wiki.hl7.org/index.php?title=Implementation_FAQ:Encryption_and_Security

HL7 MSH – Message Header. Retrieved June 1, 2017, from

<https://corepointhealth.com/resource-center/hl7-resources/hl7-msh-message-header>

HL7.org. (2017, April 19). FHIR Security. Retrieved June 1, 2017, from

<https://www.hl7.org/fhir/security.html>

Introduction to HL7 Standards. Retrieved January 31, 2017, from

<http://www.hl7.org/implement/standards/index.cfm?ref=nav>

Orion Health. Retrieved July 2, 2017, from <http://www.hl7.com/interface-engine.html>

Medical Informatics Engineering. (2016, July 6). Sample HL7 Messages. Retrieved June

1, 2017, from http://www.mieweb.com/wiki/Sample_HL7_Messages

Snell, Elizabeth. (2015, March 20). Breaking Down HIPAA: Health Data Encryption

Requirements. Retrieved March 30, 2017, from

<https://healthitsecurity.com/news/breaking-down-hipaa-health-data-encryption-requirements>