

DDS

Use Case 1: Connected Healthcare

[Return to User Experiences](#)

- ☐ **[char]Please Review**
- ☐ **[DDSFmember]Please Review**

Patient Monitoring in a Maturing Digital World

Details

Author	Matt Grubis
Title	Chief Engineer
Organization	GE Healthcare
Date	29 Apr 2020
Time	33 Minutes
Presentation	Connected Healthcare - DDSF BrightTalk
Docment	DDS in Patient Monitoring

Patient Monitoring in a Maturing Digital World

Abstract

A patient monitoring system typically consists hundreds, sometimes thousands of interconnected physiological acquisition devices (bedside monitors), real-time clinical viewers with [command and control](#) of the ecosystem (central stations), and event annunciation (sounding clinical and system alarms) to a plurality of locations, based on caregiver workflows. The [Data Distribution Service \(DDS\) middleware's Data-Centric](#), network-based state management is fundamental to the reliability, robustness and distributed operation of a system designed to make life-critical decisions and provide clinical information for further care.

Patient Monitoring as IoMT

Background

[Return to the top](#)

It is important to understand the motivation behind patient monitoring. It is not as an academic exercise or a nice-to-have but necessity for helping clinicians make the most informed decisions during life critical treatments: in emergency departments; [Intensive Care Unit \(ICU\)](#); in operating rooms; during anesthesia recovery; and, through out the birthing process for mothers and neonatals.

During patient monitor medical devices are attached to patients. The devices measure and collect vital signs and physiological states about patients in real-time.

The monitors used in Patient monitoring systems are a perfect examples of [Internet of Medical Things \(IoMT\)](#). The monitors measure and collect vital signs and physiological states about patients in real-time using wave forms and other digital formats of data. These devices by nature reside on the edge referred to in [Edge Computing](#).

The following are the basic constraints that apply to patient monitoring systems:

- Accessible all the time
- Must be reliable
- Must be unobtrusive
- Allow for patient mobility
- Support clinical flexibility
- Tolerate the use of high fidelity data
- Facilitate cyber security
- Interoperate with many devices
- Tailorable to Personal and Patient
- Participate in workflow Oriented
- Integrate with data analytics such as GE and Partner Analytics

The monitors must be well defined [endpoints](#) within an overall seamless integrated patient monitoring ecosystem.

The current GE ecosystem has three main components:

- **Patient-Centered Monitoring** - with devices attached directly to the patient
- **Flexible and Scalable Infrastructure** - with edge computing for quicker processing
- **Provider-Centered Insights** - with large datasets about the patients, their care and their history as well as analytics and [Artificial intelligence \(AI\)](#)




Figure 1: [Internet of Medical Things \(IoMT\)](#), © General Electric

An architectural decision to use the [Data Distribution Service \(DDS\)](#) is based on the basic constraints of the patient monitoring system and the understanding that the infrastructure is life critical. DDS is a foundational component that provided the support for well defined endpoints, real-time processing, redundancy, reliability and [availability](#) and allowed for the leveraging of exiting IoMT technology.

Use Case

[Return to the top](#)

Story

A man was playing racquet ball and hurts his leg, He goes to Emergency Room (ER) where he is given something to manage the pain. He is provided a self administering plunger that allows him to push a button to get a dose of morphine or other opioid. However, each individual's body reacts differently to the drugs based on their metabolism and absorption. Sometimes, an individual's respiration becomes depressed causing their heart rate to drop and the oxygenation in the lungs to decrease leading to a dangerous situation. The busy caregivers do not notice a problem until a critical event such as a heart attack occurs. Usually, the critical event just does not just "happen", but there are a series of indicators foretelling of the critical event. In other words, it is not that the patients conditions suddenly deteriorate, it is more a situation that caregivers are unaware of the deteriorating conditions until a critical event.

Data

[Return to the top](#)

- Globally, eight patients per minute die within 30 days of surgery each year. (Lancet2019)
- For many patients there is ample time prior to cardiac arrest to provide potentially life-saving interventions (CHEST2012)
- Sepsis interventions are time-sensitive and must be instituted early to achieve better outcomes. (J Intensive & CritCare 2016)
- Most post-operative deaths occur in the wards. (Lancet 2012)
- In one study, over half of the adult in-hospital cardiac arrests occurred in the general wards. (Resuscitation 2014)

USE CASE

Early detection of patient deterioration to optimize patient safety is critically important

- Globally, eight patients per minute die within 30 days of surgery each year. (Lancet 2019)
- For many patients there is ample time prior to cardiac arrest to provide potentially life-saving interventions. (CHEST 2012)
- Sepsis interventions are time-sensitive and must be instituted early to achieve better outcomes. (J Intensive & Crit Care 2016)
- Most post-operative deaths occur on the wards. (Lancet 2012)
- In one study, over half of adult in-hospital cardiac arrests occurred on the general wards. (Resuscitation 2014)

GE

Figure 2: Early detection of patient deterioration to optimize patient safety is critically important, © General Electric

Protecting the Ward Patients

[Return to the top](#)

A connected healthcare system needs to plan for patient monitoring with the idea of:

- **real-time** acquisition of data
- real-time processing of data
- analysis of the information, and
- delivering the information to clinician as quickly as possible

In fact, in patient monitoring there are critical events that occur to the patient, such as heart fibrillation basically where the heart shakes in the chest, atrial or ventricular fibrillation where a heart could stop. These are life critical events. Therefore, the system needs to assure that communication system is detecting the life critical event and to deliver information about the event to the clinician within 5 to 10 seconds. Most of the critical time is consumed just detecting the event is occurring.

In summery, the **goal** for the connected healthcare system is to be proactive of the life critical events. Early detection is the key.

- Prompt recognition and treatment of the complication is a critical, actionable point during a patient's postoperative course and can profoundly impact outcomes. (Anesthesiology 2019)
- Earlier identification and treatment of threatening conditions lead to lower mortality rates. (Resuscitation 2019)
- Each hour of delay in admission of a patient to the ICU was associated with a 1.5% increase in risk of death ... and a 1% increase in hospital mortality. (Anesthesiology 2018)

USE CASE: Protecting ward patients

Early detection is key

- Prompt recognition and treatment of complications is a critical, actionable point during a patient's postoperative course and can profoundly impact outcomes. (Anesthesiology 2019)
- Earlier identification and treatment of threatening conditions lead to lower mortality rates. (Resuscitation 2019)
- Each hour of delay in admission of a patient to the ICU was associated with a 1.5% increase in risk of death... and a 1% increase in hospital mortality. (Anaesthesiology 2018)




Figure 3: Protecting Ward Patients - Early Detection is Key, © General Electric

Patient Monitoring Overview

[Return to the top](#)

Patient monitoring is more than a device connected to a patient that acquires physiological information and processes that information generating alarms. A **Patient Monitoring system** includes many, often hundreds and sometimes thousands of devices connected to patients geographically dispersed across a hospital campus. Data and processed information from each of these devices are communicated across the hospital ecosystem and are delivered to a variety of data sinks.

The monitoring includes:

- Patient Monitoring Devices, Off-site “War Room” monitors
- Electronic medical records (EMRs)
- Mobile viewing
- Diagnostic algorithms
- Research

On the right of the circular floor plan which for this discussion is called a hospital ward. For this example, it will be referred to as Five East. Around the outside, there are patient rooms with a dozen

to 30 rooms. In the middle is a desk where some clinicians can sit watching the patients and the local monitors. Inside of each of the patient rooms there are patient monitors. In this case the ward (i.e., Five East) can be thought of as an ecosystem. But in large hospitals there may be a Five West. And there might be floors below and above these wards. Look at some hospitals like [Ohio State University Wexner Medical Center](#), [Michigan Medicine](#), [Brigham and Women's](#), [Partners Healthcare](#), [University of California San Francisco \(UCSF\) Health](#) with extremely large care areas with hundreds if not thousands of patient monitors working at the same time.

On the graphic:

[Return to the top](#)

- **[yellow dots]** Each room has an individual patient monitors.
- **[dark blue dots]** There is also central station where the clinicians can view patients data in real-time from all those patient's monitors. Even though a monitor in a patient's room is making noise or indicating some abnormality, the central station is showing the same data also. The central stations can show from 8 to 16 patients at a time allowing the clinicians to have a good overview.
- **[light blue dots]** In addition, there is a review station which allows doctors or other clinicians to review what happened before, during, and after some critical event occurred with a patient by playing back the event data stored in a log.
- **[gray dots]** To improve patient care, there are convenient viewing stations in the hallways allowing the nurses to quickly assess a patient's status without having to go into the room or without having to return to the central station.
- In addition to the ward's ecosystem, the data is also being sent to a central Electronic Medical Record (ERM) and central [servers](#) allowing mobile consults on a phone or a tablet or the doctor is doing a mobile consult from a remote location.
- The data is also sent to a cloud server where data analytics are larger analytics or population management.
- The data can also be made available to researchers that are developing algorithms to detect patient states and conditions.
- Finally, in another room or even off-site locations there is a Clinical War Room where there are humans monitoring potentially hundreds of patients providing oversight as an insurance that a clinician might have missed something. Many of the monitors in the War Rooms get to know a patient by observing their monitoring data day in and day out of the course of the patient's stay.

To summarize, this ecosystem is big, thousands of devices producing data, many devices consuming data, and highly distributed across the ecosystem and they are all connected.




Figure 4: Patient Monitoring, © General Electric

Patient Monitoring Patient State

[Return to the top](#)

The patient monitoring is not only conveying and monitoring all this information, but it also has to manage the patients' state:

- Admissions
- Discharges
- Adding and removing devices as the patients acuity changes
- Changing alarm limits
- Silencing alarms
- Reactivating alarms
- Pausing alarms for showers etc.
- Reactivating alarms
- Movements through a care area from one room to another for certain procedures
- Movements between care areas such as to the emergency room, operating room, post anesthetic care unit to critical care unit

As the patients are moving through the ecosystem the continuity managing these workflows there is a longitudinal aspect of the data as the patient moves through the care areas of the hospital.

The bottom line, the availability and [fault tolerance](#) of the system. This system has to run 24/7. It can not fail. The system has to built that once it is turned on, and thousands of patients are connected, we never turn it off. It has to keep running, even if there are failures. It has to keep running if if parts of it are disconnected.

Table 1: Patient Monitoring Requirements © General Electric.

Viewing of Physiological Data	Real-time , fault-tolerant display of waveforms, measurements and trended data.
-------------------------------	---

	Includes viewing of historical patient data, analytics and reports.
Alarming	Must sound an alarm within 10 seconds from the onset of the physiological condition.
	Failure to alarm is considered a critical failure.
Patient Management	Admit, move, discharge patients. Associate and dissociate acquisition devices.
	Configuration of patient (implanted devices, state in treatment, age-based monitoring or therapy)
Alarm Management	Clinical alarms must be delivered to all intended destinations which confirmation.
	Alarms states can be modified based on clinical conditions and algorithms.
Device and System Management	Configuration and management of all devices in ecosystem.
	Configuration and management of system. (beds, units, time, ADT, HL7, authentication, etc.)
Analytical Processing	Execution of life-critical and non-life-critical physiological algorithms.
	Results affect alarming and alarm processing throughout ecosystem.
Availability, Fault Tolerance, State Consistency	Must continue to monitor patients including in-unit central viewing, alarming and patient configuration changes if there is an infrastructure or edge-compute failure/disconnection.
	Recovery from failure should be automatic and the must perform safely in failure and during recovery.
	Must be designed with "BASE" principles (Basic Availability, Soft-state, Eventual consistency)

Patient Monitoring

Viewing of Physiological Data	- Real-time, fault-tolerant display of waveforms, measurements and trended data. - Includes viewing of historical patient data, analytics and reports.
Alarming	- Must sound an alarm within 10 seconds from the onset of the physiological condition. - Failure to alarm is considered a critical failure.
Patient Management	- Admit, move, discharge patients. Associate and dissociate acquisition devices. - Configuration of patient (implanted devices, state in treatment, age-based monitoring or therapy).
Alarm Management	- Clinical alarms must be delivered to all intended destinations with confirmation. - Alarms states can be modified based on clinical conditions and algorithms.
Device and System Management	- Configuration and management of all devices in ecosystem. - Configuration and management of system. (beds, units, time, ADT, HL7, authentication, etc.)
Analytical Processing	- Execution of life-critical and non-life-critical physiological algorithms. - Results affect alarming and alarm processing throughout ecosystem.
Availability, Fault Tolerance, State Consistency	- Must continue to monitor patients including in-unit central viewing, alarming and patient configuration changes if there is an infrastructure or edge-compute failure/disconnection. - Recovery from failure should be automatic and the must perform safely in failure and during recovery. - Must be designed with "BASE" principles (Basic Availability, Soft-state, Eventual consistency)

Figure 5: Patient Monitoring Requirements, © General Electric

The Grand Challenge

[Return to the top](#)

The Grand Challenge

The **Grand Challenge** of a Patient Monitoring Ecosystem is that the state of the ecosystem is **massively distributed**, and **always** expected to be in an *inconsistent* state, desiring a state of consistency.

Why is *inconsistency* expected?:

Consider an **average** hospital with 200 patient monitors, 10 clinical units, and 40 central stations.

It is **expected** that if a 1) Monitor becomes disconnected during transport 2) An entire Unit disconnects from the rest of the hospital 3) A building is cut off from the Edge Platform 4) The Edge Platform fails;

The state of the device, the unit, the building can be **safely** changed while *disconnected*

Safe and expected state *reconciliation*

Allowing disconnected state changes is simple, but as the elements of the system reconnect, the patient monitoring ecosystem must **reconcile** state safely, and reflect the clinical users' *intentions* based on their practice of caregiving.



A patient monitoring ecosystem is an eventually consistent system that is never truly consistent.

Figure 6: The Grand Challenge Description, © General Electric

The following helps describe what is referred to as the “Grand Challenge” of Patient Monitoring.

There is a dynamic nature to the systems

- Patients get disconnected, they enter and leave coverage areas with devices
- Network switch failures
- power outages
- fires in a building
- construction locally within the hospital but also in the streets, for example, a fiber optics line being cut, etc.

We have to build resilience into the system to adapt to the dynamic ever changing system and environment. So, while all the changes and interruptions are occurring the system is constantly trying to reconcile itself to reflect the real world. So, when the system does come back together, then the data and the system components must reconcile system to a clinically safe state.

Using **Data Distribution Service (DDS)** to maintain the state in the network and devices, allows the system to do safe and predictable reconsideration when parts of the system fail. As an example, of all the possible failures that could and do occur, you can not rely on always being in communication with the edge compute or the cloud. The system needs to back off components and expectations and continue to operate. The system needs to recover and reconcile.

- The Grand Challenge of a Patient Monitoring Ecosystem is that the state of the ecosystem is
 - **massively distributed**, and
 - **always** expected to be in an inconsistent state, desiring a state of consistency.
- **How can you design a system to be inconsistent?**
 - Because it is recognized that parts of the system fail, it is to be expected
- **Why is inconsistency expected?**
 - Consider an average hospital with

- 200 patient monitors
- 10 clinical units , and
- 40 central stations
- It is **expected** that if a
 - Monitor becomes disconnected during transport
 - An entire Unit disconnects from the rest of the hospital
 - A building is cut off from the Edge Platform
 - The Edge Platform fails
- The **state** of
 - the **device**,
 - the **unit**,
 - the **building** can be safely changed while disconnected
- **Safe** and expected state **reconciliation**
 - Allowing disconnected state changes is simple, but as the elements of the system reconnect, the patient monitoring ecosystem must
 - **reconcile** state safely, and
 - **reflect** the clinical users' *intentions* based on their practice of care giving

Graphically, the Grand Challenge is represented in the following diagram of a small system. All the components on the diagram are representative of the components with the system,. Any one of these could fail or the connections could fail. You can not be



Figure 7: The Grand Challenge, Big Perspective, © General Electric

Why Data Distribution Service (DDS) Standard

[Return to the top](#) GE is using DDS in some of the systems that were built but you don't see it , its under the hood so to speak. But as we move forward, we are continuing to expand our footprint using DDS to convey the information back and forth across the whole system.

Data Distribution Service (DDS) standard in Patient Monitoring specifically designed for real-time, mission-critical applications to manage data-centric states across decentralized systems, in a scalable and secure manner. The reason we are using DDS is:

- DDS is not “cloud centric”. If you have a server, you are dependent on the server and consequently any server failures have considered during the design. These kinds of failures are not acceptable during life critical patient monitoring.
- Data state are exchanged at “wire-speed” using **Quality of Service (QoS)** control. We prioritize some traffic over traffic. We are prioritizing real-time event traffic versus historic back logs. For example, if a device had been disconnected for four hours, the backlog data is important and needs to be back filled but not interfere with the real-time monitoring data. This prioritization of the data delivery is done using a QoS scheme across all the different topics and domains in the system. Naturally, it is optimized for the real-time aspects of patient monitoring.
- There are no data-mastership concerns, just distributed state reconciliation. In the past, before we used DDS, some data was stored in the device, some in the central station, and some on the servers. This made scaling and extending the system difficult because each data repository had different data and data-mastership issues.
- Efficient communication leveraging DDS built-in “type system” (**DDS-XTYPES**). Patient Monitoring data is complex because human physiology is complex. Some example of the data devices collect are:
 - electrical impulses from the heart with ECG
 - saturation levels oxygen (i.e., oxygenation) in the blood
 - brain waves,
 - exhaled respiratory gasses such as CO2

There are so many measurements that can be collected from a body, it results in very complex **Data Model (DM)**. Using verbose, un-managed data models like those in HL7 is not efficient and can lead to errors of translation. So, representing the data the data using DDS X-TYPES, the

data is ultimately represented as binary data. It is efficient from at the machine level. The data can be sent across the wire in binary format, but is still fully documented and described with XTYPES. So each station shares the same current data model regardless of its computer architecture (i.e., big versus [little endian](#), [16-Bit](#), [32-Bit](#) or [64-Bit](#) processors, etc.

- On the wire compatibility with dynamic data model dynamic changes and evolution. This provides on-wire compatibility which is really important as the data model changes. The data model not only changes through the current care workflow on different machines, but also aa the system is extended and upgraded in the future. For example, providing upward compatibility as parameters are modified or added going forward, or even using new parameters we have not even invented yet.
- Content aware filtering to drive further efficiency - A lot of our devices are wireless, battery operated and using [wan](#) connectivity. Sending un-needd or unwanted data to and from these devices can shorten their usefulness or create the the equivalent of a Denial of Service (DoS) attack. Therefore, we have to make sure for efficiency and security that the data transfers are efficient and minimal. In other words, we are not sending data to places in the network that do not need or want the data. So, we rely on the built-in [Data Distribution Service \(DDS\)](#) content filtering at both the source and the destination. So we are sending the correct data to the data sinks that require that data and nothing extra.
- Fully secure endpoint authentication and [encryption](#) built-in - In healthcare, security is paramount. Making sure that we have military grade security, there is encryption, authentication, authorization and that different nodes have different rights to different topics and different ways of exerting Command and Control

(C2) over a patients state (i.e., admitting, discharging, silencing alarms, changing the way an algorithm behaves.) These types of features must be at the foundation of the system, otherwise the entire system at risk. In our system, we are using DDS-Secure to help with these issues.

- Transport independent, the best transport is used for the communication link - Some of the links are wireless LAN, some are over wired high-speed networks, and others are over point-to-point “body area” wireless networks. Not everything is Ethernet or Internet Protocol (IP), therefore, we use a service which not tied to a specific transport allowing the right transport to be selected for the right purpose. This applies all the way along the communication pathway from the edge to the servers.
- Software defined domains and topics - Help shape the distribution of data, the virtual communication of which endpoints communicate with each other, also adding to the design architectures and the security and efficiency of the system.

DDS in Patient Monitoring

The Data Distribution Service standard was specifically designed for real-time, mission-critical applications to manage data-centric states across decentralized systems, in a scalable and secure manner.

- DDS is not “cloud centric”
- Data state is exchanged at “wire-speed” with Quality of Service control
- There are no data-mastership concerns, just distributed state reconciliation
- Efficient communication leveraging DDS built-in “type system” (DDS-XTYPES)
- On the wire compatibility with dynamic data model dynamic changes and evolution
- Content aware filtering to drive further efficiency
- Fully secure endpoint authentication and encryption built in
- Transport independent, the best transport is used for the communication link
- Software defined domains and topics



Figure 8: DDS Features that help achieve the Grand Challenge of Patient Monitoring, © General Electric

From:
<https://omgwiki.org/dds/> - DDS Foundation Wiki
 Permanent link:
https://omgwiki.org/dds/doku.php?id=dds:public:guidebook:03_user:02_health&rev=1626292082
 Last update: 2021/07/14 15:48

