Health Analytics as a Service with Artemis Cloud: Service Availability*

Carolyn McGregor *Senior Member, IEEE*, Catherine Inibhunu, Jonah Glass, Ian Doyle, Aaron Gates, John Madill, J. Edward Pugh

Abstract—**Critical care units internationally contain medical devices that generate Big Data in the form of high speed physiological data streams. Great opportunities exist for systemic and reliable approaches for the analysis of high speed physiological data for clinical decision support. This paper presents the instantiation of a Big Data analytics based Health Analytics as-a-Service model. The availability results of the deployment of two instances of Artemis Cloud to support two neonatal ICUs (NICUs) in Ontario Canada are presented.**

Clinical Relevance— **This research demonstrates the robustness of Artemis provisioned as a cloud-based service to neonatologists for Big Data based Clinical Decision Support for improved clinical decision making in Neonatal Intensive Care.**

I. Introduction

Critical care units internationally provide acute care for patients in critical conditions involving interdisciplinary teams of healthcare workers. For over five decades, medical devices have been generating and displaying streams of patient data at high speed with numeric data such as derived heart rates that are calculated and displayed each second (1Hz) and waveform data such as electrocardiograms producing data at rates of 512 or 1000 readings a second. Data speeds of this form are necessary in critical care as they are measuring events in human physiology such as the frequency of the beating heart and degree of oxygen perfusion in the body to support the health of vital organs such as the heart, brain and lungs. The output mechanism and data formats vary across medical devices. Many devices have proprietary standards. This form of data has been characterized as Big Data due to its volume, velocity and variety. Solutions that enable the real-time effective use of this raw high speed physiological data for the detection of known and discovery of new pathophysiologies for a range of conditions remain a challenge [1], [2].

Neonatal intensive care units (NICUs) are critical care units that care for premature and ill term infants. Premature infants within a NICU can be as young as 22 weeks gestation and can be cared for within an NICU for 6 months or longer depending on their medical conditions. Detection of a range of neonatal conditions remain a challenge and some conditions such as late onset neonatal sepsis, can cause death within hours of detection of clinical decline at the bedside [3]. As a result, great opportunities exist for systemic and reliable approaches for the analysis of high speed physiological data for clinical decision support. However, to date, deployment of clinical decision support has largely been limited to the addition of modules within bedside monitors rather than a systemic platform based approach.

Stream computing is a newer form of computing architecture that enables the processing of high velocity Big Data for real-time/low latency derived analytics. Stream computing addresses a fundamental assumption that the value of the analysis of streams of data is in its freshness and real-time nature [4]. This paradigm is well suited to the problem of real-time analysis of physiological data for new approaches to clinical decision support.

Artemis Cloud is a Big Data analytics platform offered as a Health Analytics as-a-Service (HAaaS) through the cloud. Artemis Cloud supports both real-time and retrospective analysis of high speed physiological data along with other clinical data to support clinical decision support and clinical research respectively in critical care. The real-time component of Artemis is provisioned through stream computing. This HAaaS architecture represents a minimal footprint approach within the healthcare organization and the analytics platform components exist outside the organization.

In order to provide a viable solution for the real-time clinical decision support component, clinical decision support systems such as Artemis Cloud must demonstrate an ability to provide a high degree of availability at the bedside. This paper presents the instantiation of a Big Data analytics based Health Analytics as a Service model of care that was previously modelled in [3], [5]. The details relating to two Artemis Cloud instances in a pre-production context are presented to gather evidence of the reliability of Artemis Cloud as a Health Analytics as-a-Service model. The availability results of the deployment of two instances of Artemis Cloud to support two NICUs in Ontario, Canada are presented.

II. Related Work

The first prototype architecture for Artemis was proposed through a collaboration between McGregor and IBM as a platform to support real-time clinical decision support in neonatal intensive care along with supporting retrospective medical research studies [6]. That version of Artemis was deployed using IBMs InfoSphere Streams for the real-time

*Research supported by FedDev Ontario as part of the Health Ecosphere Innovation Pipeline Project, the Canada Research Chairs program (#950-225945); the Canadian Foundation for Innovation (#203427); and McGregor's Research Excellence Chair in Health Informatics.

C. McGregor is with Joint Research Centre in AI for Health and Wellness, Ontario Tech University, Oshawa, Canada and University of Technology Sydney, Australia (e-mail: c.mcgregor@ieee.org).

C. Inibhunu, J. Glass, I. Doyle, A. Gates, J. Madill are with the Joint Research Centre in AI for Health and Wellness, Ontario Tech University, Oshawa, Canada.

J. E. Pugh is with McMaster Children's Hospital, Hamilton, Canada and Ontario Tech University.

analytics of the physiological data streams and DB2 for the data storage with each operating on their own dedicated laptops. This demonstrated the potential of the InfoSphere Streams software and was a key driver in the release of InfoSphere Streams by IBM as a product. A Stealth CPU device ran CapsuleTech software as the Data Acquisition interface to Philips Intellivue monitors. CapsuleTech serial to ethernet converter devices were installed to receive a subset of the physiological data streams generated by the Intellivue devices from up to four bed spaces in the collaborating NICU.

An initial design for Artemis Cloud was outlined in [7] where a series of web services were proposed to interface with Artemis to enable the Data Integration, Data Analytics and Data Visualization components to operate remotely from the hospital. This was demonstrated in a case study with a second collaborating NICU [1], [2]. However, that architecture although maintaining a high level of service availability was not designed as a robust deployable offering.

An extended Artemis Cloud architecture enabling it to operate as a Health Analytics as a Service for clinical use was previously designed and modelled in [3], [5].

Wang [8] proposes a best practice big data analytics architecture that supports the capture, transformation and consumption of big data in healthcare. It contains five key architectural components namely: the data layer; the data aggregation layer; the analytics layer; the information exploration layer and the data governance layer. However, that model does not consider the provision of the data analytics architecture in the cloud and does not consider the evaluation of the availability assessment of these forms of architectures.

Baby and Ravikumar [9] assessed architectures of Big Data technologies from the perspective of speed, online processing, scalability and reliability and found the design of Artemis to have the greatest potential and highest score across each category. A key challenge of healthcare Big Data analytics platform architectures proposed is their potential and measured ability for service availability, reliability and scalability.

In their review of the state of the art, challenges and opportunities of the application of machine learning for streaming data Gomes et al [10] note that there "are many challenges to be addressed before existing methods can be efficiently applied to real-world problems" and include evaluation of the availability and processing of the data as some of these challenges. Kolajo et al [4] note there are challenges arising from the nature of processing Big Data relating to availability associated with scalability, fault-tolerance, timeliness, stability and high throughput. Hence the importance of reporting on the availability of the real-time component of a Big Data platform such as Artemis Cloud.

III. Methods

The Artemis Cloud Health Analytics as a Service architecture is presented in Figure 1.

A. Data Collection

Within critical care, medical devices are connected to patients via a range of sensors to gather data about their current health state. These devices transform that signal into high frequency data streams that are usually displayed as numbers or waveforms on monitors at the bed space. Artemis Cloud connects to existing medical devices to acquire and further utilize this data that has been collected.

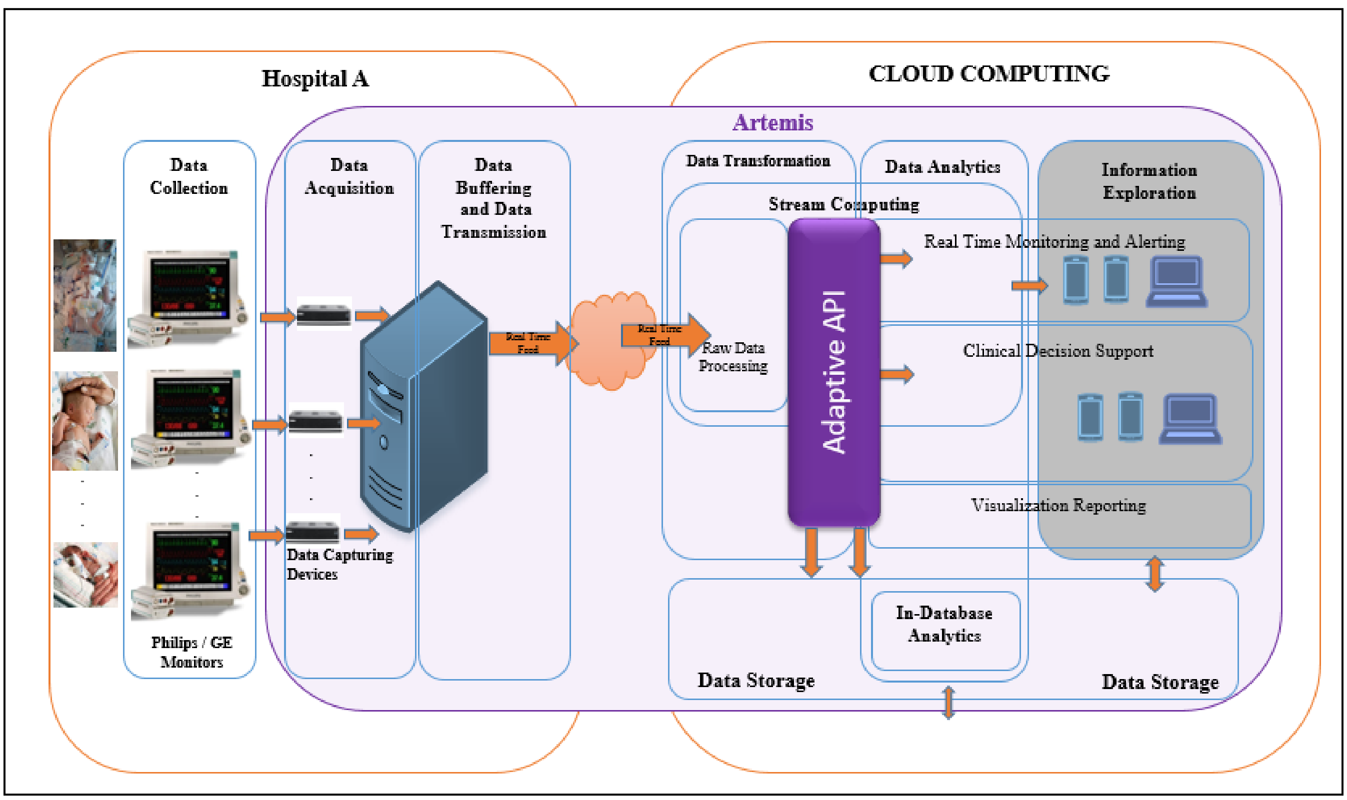

Figure 1. Artemis Cloud [11]

B. Data Acquisition

Artemis Cloud acquires data from existing bedside medical devices enabling use of Artemis Cloud with minimal impact.

C. Data Buffering and Transmission

The process of data acquisition is performed by Data Buffering and Transmission software whose function is to transmit data onto the Adaptive API of Artemis Cloud and to buffer the data when transmission functions are not available due to outages on the network or the Adaptive API.

D. Data Transformation

As noted in [11] the Artemis Cloud API allows for the storage of the raw data within the database management system. In addition, it supports multiple data ingestion workflows by preparing the data for consumption by other services such as multiple analytics algorithms.

E. Data Analytics

The Data Analytics component enables the deployment of several data analytics modules that can consume the data made available by the Adaptive API. This enables summary analytics reflecting information about the patient state from the perspective of a range of medical conditions to be generated.

F. Data Storage

The raw data along with the resulting temporal abstraction data are stored within the Data Storage component.

G. Information Exploration/Visualization

Artemis Cloud enables exploration and visualization of information through the information exploration and visualization components.

IV. Case Study

The Artemis Cloud architecture presented in Figure 1 was deployed within the NICUs of two hospitals in Ontario, Canada as a flagship project within the Health Ecosphere Innovation Project funded by FedDev Ontario. It is deployed in the Level 3 portion of the NICU in McMaster Children Hospital (MCH). The second deployment is in the Level 2b NICU in Southlake Regional Health Centre (SRHC).

MCH is a children's hospital located in Hamilton, Ontario, Canada. It has the largest NICU in Ontario. There are 48 bed spaces within the Level 3 portion of the NICU, together with 3 bed spaces within the delivery suite.

SRHC is a regional hospital located on the outskirts of the Greater Toronto Area in Ontario. The Level 2b NICU at SRHC contains 17 bed spaces. Infants within the Level 2b NICU have a lower level of acuity than those in the Level 3 NICUs.

These two NICUs were strategically chosen for the Artemis Cloud deployments based on the different medical monitors they utilized in the NICU and the different levels of acuity.

Privacy Impact Assessments together with Threat and Risk Assessments documentation was completed at both MCH and SRHC and was approved by each hospital's IT Departments prior to the installation of the Artemis in hospital components. An architecture consisting of two separate Artemis Cloud deployments, one for each hospital, rather than one Artemis Cloud deployment with logical separation of hospital data was chosen by hospitals due to security and privacy concerns of healthcare data. The cloud environment for each Artemis Cloud instance for each hospital was provisioned by the Centre for Advanced Computing (CAC), Queen's University as a node within the Compute Ontario advanced research computing infrastructure in Ontario. These space allocations were provisioned as 'bear metal' with no virtualized infrastructure and the Artemis Cloud deployments were built and maintained by the Ontario Tech University IT Services division. The CAC provided 24/7 support for the availability of the Artemis Cloud infrastructures and maintenance windows on those infrastructures were arranged as planned outage windows.

Ethics approval for the Artemis Cloud deployment research studies at MCH and SRHC were received from MCH (#3859-D), SRHC (#0022-1718) and the University of Ontario Institute of Technology now known as Ontario Tech University (#14736 for MCH and #14667 for SRHC).

A. Data Collection

The Artemis Cloud deployment at MCH NICU collects data from the existing Philips Intellivue MP70 monitors located at each NICU and delivery suite bed space. At SRHC, the deployment collects data from the GE Dash monitors at each bed space in the NICU.

B. Data Acquisition

The Phillips Intellivue devices in the NICU at MCH have an RJ45 serial port that was used for connection to Artemis. These devices have an ethernet port also that was being used by the Philips central monitoring software. To enable the conversion of the data output from the serial port to ethernet, a Digi serial to ethernet conversion device (Digi DC-ES-4SB-SW) is located at each bed space. A four serial port to one ethernet serial to ethernet conversion device was chosen for each bed space to allow for future expansion and connection of other bedside devices that have serial outputs. This approach enabled us to have a separate network connection to the medical device than that which continues to be used by the Philips central monitoring software.

For SRHC, due to challenges with direct device connectivity as a result of proprietary protocols, we chose to install the GE Carescape central monitoring module and connect each monitor to that module and then acquire the data from GE Carescape.

We utilized the off-the-shelf Vines software, from True Process (owned Baxter), to perform the data acquisition function. The Vines software has interfacing software to several medical devices. However, the list price of the Vines software was over $1000 per bed space plus ongoing maintenance. The Vines software performs three key functions as part of the data acquisition process namely, acquiring the data, translating the medical record number (MRN) to an alternate deidentified number for transmission outside the hospital and providing the ability to buffer the data when either the network or the Artemis API is offline.

At both MCH and SRHC a dedicated server was installed to house the Vines software to perform data acquisition. At each hospital the Vines software establishes a secure connection to its Artemis Cloud instance at Queens via Open VPN. The server is "security hardened" and is maintained from a patching and vulnerability through use of patching software. This creates short windows each month of planned downtime for security patching. Monitoring the health of the server housing Vines at each hospital was performed by an Artemis Cloud operational team within the Ontario Tech University IT Services division.

At MCH, the patient MRN is entered into the Philips Intellivue device as part of routine care within the NICU.

At SRCH nursing staff, as part of the study, entered the MRN for the patient into the medical device as this was not previously part of routine care.

C. *Data Buffering and Transmission*

Data is transmitted from Vines to the API within each Artemis Cloud instance located at the Centre for Advanced Computing (CAC), Queen's University through the ORION Network. ORION is the Ontario Research Innovation Optical Network that and is the province of Ontario's only provincial network for research, healthcare and educational institutions (www.orion.on.ca). The ORION IP at each hospital utilized by Vines is only routable to the CAC at Queen's University.

The Open VPN connection between MCH and the Artemis Cloud instance for MCH at the CAC is 1GB. The Open VPN connection between SRCH and the Artemis Cloud instance for SRHC at the CAC is 100MB.

D. *Data Transformation*

The Artemis Cloud API is implemented within IBM Infosphere Streams as a Streams graph.

E. *Data Analytics*

We have implemented 9 Streams graphs as analytics algorithms consumers that subscribe to the adaptive API which then process the data streams and sends the resulting analytics for storage and/or presentation. These analytics algorithms perform a range of temporal abstraction functions on different sets of raw data streams to enable a range of clinical decision support functions. Temporal abstraction analyses windows of data streams to derive meaning from the raw data stream. Temporal abstractions can be simple abstracts based on one data stream or multiple data streams or hierarchical sets of temporal abstractions. For each of the two instances of Artemis Cloud, the 9 analytical algorithms that have been deployed as consumers of the implemented API. The algorithms perform temporal abstraction functions for the clinical conditions of:

- Late Onset of Neonatal Sepsis [12];
- Retinopathy of Prematurity (ROP) [13].

F. *Data Storage*

The raw data along with the resulting temporal abstraction data are stored within the Data Storage component provisioned in IBM DB2. DB2 was chosen as the data is highly structured, and a high level of service reliability and availability was required.

G. *Information Exploration/Visualization*

Within this first stage of the Artemis deployment, the focus was on demonstrating the reliability and availability of the Artemis Cloud platforms in real-time to demonstrate the robustness of the architecture. As a result, no real-time visualization was provided at the bedside for clinical decision support. However, the results of the retrospective analysis of the performance of the algorithms in relation to a set of neonatal conditions is the subject of other publications currently and will be reported in the future.

V. RESULTS

In each hour there are up to 68 patients across the two NICUs online and data acquired from these patients is transmitted to Artemis Cloud in real-time.

The Artemis Platform availability at MCH and SRHC was evaluated by reviewing all incidents within the Information Technology Infrastructure Library Incident Log and analyzing the downtime, reason for the downtime, the components and number of beds affected, and data recovery time.

A. *McMaster Children's Hospital*

In one hour approximately three million data points are processed for each infant by Vines and forwarded to the Adaptive API component for processing as JSON messages. When all 51 bed spaces are occupied at MCH for example, then approximately 153 million data points per hour are processed by each of these components within the MCH Artemis Cloud instance.

The Artemis Platform deployed at McMaster Children's Hospital was evaluated between the period of March 1, 2018 and September 30, 2018. A period of 213 days, or 306,720 minutes. During the period, 38 incidents that affected system availability were listed in the incident log.

A summary of the resulting availability metrics are below

TABLE I. ARTEMIS CLOUD AVAILABILITY MCMASTER CHILDRENS HOSPITAL

McMaster Children's Hospital (MCH)		
Metrics	***Result***	***%***
Total System Minutes:	306,720	
Total Beds	51	
Total Incidents	38	
Total Minutes of Documented Unplanned Downtime (as a single system)	10,158	
Total Minutes of Documented Planned Downtime (as a single system)	126	
Availability percentage of Artemis as a singular system	296,436 / 306,720	96.6% Availability
Count of beds with availability percentages greater than 99.5%	45	88% of beds

B. *Southlake Regional Health Centre*

The Artemis Cloud Availability for SRHC is summarized in Table II.

TABLE II. ARTEMIS CLOUD AVAILABILITY SOUTHLAKE REGIONAL HEALTH CENTRE

Southlake Regional Health Centre (SRHC)		
Metrics	***Result***	***%***
Total System Minutes:	306,720	
Total Beds	17	
Total Incidents	6	
Total Minutes of Documented Unplanned Downtime (as a single system)	578	
Total Minutes of Documented Planned Downtime (as a single system)	244	
Availability percentage of Artemis as a singular system	305,898 / 306,720	99.7% Availability
Count of beds with availability percentages greater than 99.5%	17	100% of beds

VI. CONCLUSION

This paper has presented the initial deployment results of an instantiation of a Big Data analytics based Health Analytics as a Service model, Artemis Cloud. Assessing the results of the deployment from an availability perspective is important to demonstrate the viability of this form of Big Data analytics platform to support clinical decision support in critical care. Across the two deployments Artemis maintained service availability over the six month period of 96.6% and 99.7% respectively and for individual bed spaces, the number of bed spaces where availability was greater than 99.5% was 88% and 100% respectively.

ACKNOWLEDGMENT

The authors wish to thank all the infants and their families that participated in this research study in the NICUs of MCH and SRHC. We thank Pat Clifford and the IT division at SRHC for their assistance with the deployment of Artemis at SHRC. We thank HITS within Hamilton Health Sciences for their assistance with the deployment of Artemis Cloud at MCH. The authors wish to thank Geoff Travis for consenting patients for the study at McMaster Children's Hospital and David Barbosa for enrolling patients at Southlake Regional Health Centre. The authors wish to thank Pat Clifford for his support of this research and the nursing staff in the NICU at Southlake Regional Health Centre for admitting and discharging patients through the Dash monitors for the study.

REFERENCES

[1] C. McGregor, "Big Data in Neonatal Intensive Care," *Computer (Long. Beach. Calif).*, vol. 46, no. 6, pp. 54–59, Jun. 2013.

[2] C. McGregor, "Big Data in Critical Care Using Artemis," in *Signal Processing and Machine Learing for Biomedical Big Data*, E. Sejdic and T. H. Falk, Eds. CRC Press, 2018.

[3] H. Khazaei, N. Mench-Bressan, C. McGregor, and J. E. Pugh, "Health Informatics for Neonatal Intensive Care Units: An Analytical Modeling Perspective," *IEEE J. Transl. Eng. Heal. Med.*, vol. 3, 2015.

[4] T. Kolajo, O. Daramola, and A. Adebiyi, "Big data stream analysis: a systematic literature review," *Journal of Big Data*. 2019.

[5] H. Khazaei, C. McGregor, J. M. Eklund, and K. El-Khatib, "Real-Time and Retrospective Health-Analytics-as-a-Service: A Novel Framework," *JMIR Med. Informatics*, vol. 3, no. 4, p. e36, 2015.

[6] M. Blount *et al.*, "Real-time analysis for intensive care: Development and deployment of the artemis analytic system," *IEEE Eng. Med. Biol. Mag.*, vol. 29, no. 2, pp. 110–118, 2010.

[7] C. McGregor, "A Cloud Computing Framework for Real-time Rural and Remote Service of Critical Care," in *Computer-Based Medical Systems*, 2011, p. CD-ROM 6 pages.

[8] Y. Wang, L. A. Kung, and T. A. Byrd, "Big data analytics: Understanding its capabilities and potential benefits for healthcare organizations," *Technol. Forecast. Soc. Change*, 2018.

[9] K. Baby and A. Ravikumar, "Big Data : An Ultimate Solution in Health Care," vol. 106, no. 10, pp. 28–31, 2014.

[10] P. Gomes, M. Kaiseler, B. Lopes, S. Faria, C. Queiros, and M. Coimbra, "Are standard heart rate variability measures associated with the self-perception of stress of firefighters in action?," in *Proceedings of the Annual International Conference of the IEEE Engineering in Medicine and Biology Society, EMBS*, 2013.

[11] C. Inibhunu, R. Jalali, I. Doyle, A. Gates, J. Madill, and C. McGregor, "Adaptive API for Real-Time Streaming Analytics as a Service," in *41st International Engineering in Medicine and Biology Conference*, 2019, pp. 3472–3477.

[12] C. McGregor, C. Catley, and A. James, "Variability analysis with analytics applied to physiological data streams from the neonatal intensive care unit," in *Proceedings - IEEE Symposium on Computer-Based Medical Systems*, 2012.

[13] K. E. S. Fernando, C. Mcgregor, and A. G. James, "CRISP - TDM 0 for Standardized Knowledge Discovery from Physiological Data Streams : Retinopathy of Prematurity and Blood Oxygen Saturation Case Study," in *IEEE Life Sciences Conference*, 2017, p. 4 pages.