An evaluation of the insidious consequences of clinical computing infrastructure failures at a large academic medical center

Sandeep Malleshwar Napa

A thesis

submitted in partial fulfillment of the

requirements for the degree of

Master of Science

University of Washington

2019

Committee:

Mark E Whipple

Thomas H Payne

Program Authorized to Offer Degree:

Biomedical and Health Informatics

University of Washington

Abstract

An evaluation of the insidious consequences of clinical computing infrastructure failures at a large academic medical center

Sandeep Malleshwar Napa

Chair of the Supervisory Committee:

Mark Whipple, Associate Professor

Biomedical Informatics and Medical Education

Electronic Health Records (EHRs) are intended to make healthcare delivery safer, more effective and accountable. They are complex socio-technical systems that are dependent on the proper functioning of many individual components that comprise the clinical computing infrastructure (CCI), such as networking equipment, message routing systems, departmental clinical computing systems and many others. However, on occasion these CCI components fail or need maintenance, causing clinical workflow and data flow disruptions called downtimes. Considering the inherently disruptive nature of EHR downtimes, organizations typically have mitigating procedures in place. However, many other small hardware or software CCI components also fail, causing loss of EHR functionality, sometimes insidiously. To our knowledge, systematic analysis

of CCI failures has not been undertaken. In this work, a dataset of CCI failure logs gathered at one health care system was classified and categorized to shed light on the nature, diversity, frequency and user impact of such failures. ORCA (Online Record of Clinical Activity), Epic and Mindscape are EHR components. By number of records, the top 3 components that had the highest frequency of failure are: Network (393 incidents, 59.5% of which were unscheduled) the inpatient EHR (ORCA) (372 incidents, 49.5% unscheduled) the outpatient EHR (Epic) (228 incidents, 12.3% unscheduled). In terms of user impact, components that accumulated the most failures are: the inpatient EHR (ORCA) (284.8 hours among under 5 users), Cloverleaf (interface engine) (263.5 hours among under 200 users), imaging (205.8 hours among under 50 users), and network (193.9 hours among under 50 users, and 193.4 hours among under 10 users). It is interesting to note that 4 of the 5 aforementioned components affected under 50 users. For the data with user impact estimates, cumulative EHR downtime (687.3 hours) is less than cumulative downtime for small impact non-EHR CCI failures (1131.1 hours). So, it is possible that cumulatively these small-impact but more frequent CCI component failures may approach or exceed the clinical impact of EHR downtimes. Although the data used in this work have important limitations in their accuracy and completeness, this exploratory analysis is the first step towards a better understanding on how to build a safe, resilient CCI that more reliably serves the needs of patients and providers.

TABLE OF CONTENTS

INTRODUCTION

Background

Electronic health records (EHR) have become ubiquitous and are being increasingly relied upon to deliver patient care since the passage of the American Recovery and Reinvestment Act in 2009.[1] As of 2015, more than 88% of healthcare institutions have adopted an EHR.[2] As a result, many healthcare processes are increasingly accomplished via an EHR, which has become critical for the flow of clinical information, decision making, and care coordination. At the University of Washington's (UW) two main teaching hospitals, UWMC (University of Washington Medical Center) and HMC (Harborview Medical Center), each day about 1,200 imaging reports and 50,000 laboratory results are available. After review, 30,000 electronic orders and 9,000 physician notes are written. Nurses, pharmacists and therapists author another 3,500 notes.

The smooth functioning of such a large distributed system is dependent on many components, which vary depending on the organization and even within the same organization over its history. These components, could include among others: departmental systems (picture archiving and communication systems, radiology information system, laboratory information system, anatomic pathology, cardiology systems for ECG, cardiac catheterization, echocardiography, inpatient pharmacy, retail pharmacy), medication dispensing, surgical scheduling, anesthesiology, radiation oncology, interface engines, networking components and Citrix servers to deliver applications to the point of care workstations (fixed and mobile). EHR and its associated components are what can be described as Clinical Computing Infrastructure (CCI).

On occasion, CCI components fail and cause clinical workflow disruptions, called downtimes. Major EHR downtimes are a reality at most organizations. Unplanned EHR downtimes can cause significant patient safety problems because clinical operations and decision making are so dependent on their proper functioning[3-7]. Downtime planning involves a workstation running a truncated version of EHR with paper-based ordering and documentation system. This is true whether EHR is run remotely or locally. The transition to and from paper processes creates its own risks of errors of duplication, omission and the loss of time spent verifying and communicating clinical orders and results[8]. Users may not be sufficiently trained in downtime procedures, which also poses risk. EHR downtimes are obvious to clinicians because of their inherently disruptive nature. These events are usually planned and have downtime procedures to offset the clinical impact of these disruptions. Non-EHR, CCI failures, on the other hand, are also, and perhaps more, difficult to predict. Importantly, clinicians may not be even aware of the problem and may incorrectly believe that they are working with the most recent information about the patient. They may even mistakenly assume that their orders are being communicated, leaving the potential to cause patient harm. If clinicians don't see the most recent information, and if they are aware that is the case, they may postpone important decisions or seek information from other sources such as calling the radiology department or paging a colleague. If, on the other hand, they are not aware that the display they are accustomed to viewing for the most recent data, instead does not display the most recent data, they may proceed with actions that are decidedly different than actions they would take with current data. They may be unaware that some colleagues do not have access to the EHR when a portion of the network fails.

The term “failure,” is used here but sometimes the reason for component unavailability could be scheduled maintenance or upgrades that clinicians may be unaware of, despite efforts to notify them that they do not see. At UW, multiple methods are used to notify our users: banners, electronic mail, posted notifications and working group meetings. Notification strategies are not completely effective; some users are unaware of the change in information availability despite attempts to notify them.

It can also be hard to develop outage protocols for these CCI failures. Younger generations of clinicians may not be proficient at writing orders on paper. The clinical impact of scheduled and unscheduled downtime because of non-EHR CCI failure or necessary software upgrades may cumulatively exceed an EHR downtime alone. The impact is variable between organizations depending on the computing architectures used, downtime processes followed, and organizational experience in managing CCI failures.

Health information technology (HIT) can be described as a socio-technical system in which the elements of people, processes and technology interact with each other, the organization, and the external environment in which it is employed [9-10]. This understanding has proved useful to acknowledge the successes, failures and dangers in terms of patient harm of HIT implementations.

Sociotechnical Model of Clinical Computing Infrastructure

Adapted from Harrington et al. (2010), Sittig and Singh (2010) and Walker et al. (2008)

The interdependency component of the socio-technical system is important because failure of departmental systems, foundational systems, interface engines, terminal servers, network components, point of care devices and other infrastructure components can manifest as an interruption to the EHR functioning. Sometimes users may perceive they are experiencing a traditional EHR downtime but are misled, because if a terminal server (hardware or application) fails only those using the affected terminal server to connect to the EHR will face service interruption. These users may believe that the EHR is down, when in reality it is not. If a network component servicing an intensive care unit (ICU) fails, users may believe that EHR has failed. Departmental systems may not be able to send data to EHR due to a failure of interface engine, terminal server, network or other components but end users may perceive it to be an EHR failure, probably because EHR is the interface that represents the human element of data interaction.

There is some literature describing high-level descriptions of incidents [11-13], analyses of failures of specific departmental systems[14,15], analysis of human errors or workflow related errors[16] and single incident case reports[17-19]. Reports on EHR associated safety risks only describe the risks of using an EHR[20]. The HIT Hazard Manager framework by AHRQ (Agency for Healthcare Research and Quality) describes failures as 'User hardware not working or malfunctioning" and "Back-end hardware failure". This approach does not adequately address the classification of infrastructure failure[21]. The more recently published SAFER guides (Safety Assurance Factors for EHR Resilience) provide guidance on health IT safety improvements but lacks detailed analysis of infrastructure failures.[22]

A deeper understanding of the frequency, types and clinical implications of CCI failures can help us prevent and mitigate the insidious harm they have the potential to create. An analysis of such infrastructure failures has not yet been identified as important in the evolving understanding of EHRs and other clinical computing environments.

Objectives

The impact of CCI failures on the safety and efficacy of patient care is a function of many factors such as the nature of failure (scheduled or unscheduled), clinician awareness of the malfunction, number of hours of downtime, number of clinician users affected, and frequency of failure. The most important factors, however, are the criticality of the component in providing patient care, architecture of the CCI and the organizational experience around mitigation strategies. The first goal is, therefore, to create a visual model of the CCI by interviewing health technology experts at our institution. This model will illustrate dependencies and the consequences of CCI failures. The next objective is to analyze our dataset of system logs to provide estimates of infrastructure failure types and frequencies. Finally, the clinical consequences and patient safety risk of CCI failures, are explored, by interviewing clinicians.

METHODS

Study setting

The dataset used in this work is a log of clinical computing infrastructure compiled by Dr Payne. This log was created in his capacity as the Medical Director for Information Technology Services for the University of Washington care delivery system. He has responsibility for assisting the IT Services group in avoiding disruption of the CCI used by health care providers. The logs, messages and files it contains are used in day-to-day tracking of CCI problems at HMC and UWMC, both principal teaching hospitals for the University of Washington. As of 2009, HMC had 413 licensed beds, 4,432 employees, and 1,216 physicians. There were 19,424 admissions, 65,515 Emergency Department visits, and 13,455 surgery cases. As of 2010, UWMC had 450 licensed beds, 3,982 employees, 1,829 physicians and 1,000 residents and fellows. There were 19,260 admissions, 25,602 Emergency Department visits and 15,137 surgery cases. The dataset is only an approximation of the reported CCI failures at UWMC (University of Washington Medical Center) and HMC (Harborview Medical Center). Incidents from affiliated UW Neighborhood Clinics (Ballard Clinic, Bell town Clinic, Federal Way Clinic, Issaquah Clinic, Kent/Des Moines Clinic, Northgate Clinic, Olympia Clinic, Orcas Island Clinic, Ravenna Clinic, Shoreline Clinic, Smokey Point Clinic, South lake union Clinic, Woodinville Clinic) are also included.

Study Design

The study was conducted in three phases:

Phase 1: Meetings with advisors and technical team members to define the CCI architecture at HMC and UWMC

Phase 2: Analysis of the logs of system failures gathered by the Dr Payne to create a database of events with the following elements for each event.

- Date of email and time of email sent
- Date of actual downtime event in consideration
- CCI component involved
- Root cause of failure, if available
- Estimated number of affected users
- Estimated number of hours in downtime
- Whether failure was scheduled, or unanticipated.

Phase 3: Follow up interviews to get context for data and gather clinical vignettes to corroborate patient and end user impact.

Data Analysis

The data set was first extracted from compressed zip files. Not all data fields describing details of the failure were available for all incidents. The downtime periods, if unavailable were taken from the follow up email report or calculated from the initial email to time of resolution 'all clear' email. The number of users affected was a best guess estimate by the person who authored the incident report at the time the report was filed. An incident or event had one or multiple emails (one to notify users of the issue, sometimes updates, one to notify resolution) Sometimes the text files and screenshots represented knowledge the emails already had. Care was ensured to reduce redundancies while entering the information, with two rounds of data cleaning to check for redundancies, and one last round of checking when grouping component types. The data was first complied into a Microsoft Excel spreadsheet, ver 16.11.1 (Microsoft Corporation, Redmond) then visualized using Tableau Windows desktop edition, ver 2018.1.7 (Tableau Software, Seattle)

Inclusion/Exclusion criteria

All systems that had a direct clinical impact, business systems that had an impact on clinical operations, and analytics systems were included. Systems that didn't have a clinical impact, systems that the experts deemed insignificant impact, training systems, workforce management systems, facilities management systems, Non-IT systems, Network monitoring systems, and Performance management systems were excluded. Our data set is a collection of 5019 files comprised of messages and other files as an outcome of tracking CCI failures between August 20, 2001 and January 21, 2018, a span of 16 years and 5 months. This includes 4805 emails in eml format (95.7%), 136 rich text files in rtf format (.027%), 5 rich text files in rtfd format (.001%), 16 plain text files in txt format (.003%), 16 PDF files (.003%) and 13-word documents in doc format (.003%). Some logs were captured as screengrabs and image files. We have 16 images in jpg format (.003%), 10 images in png format (.002%) and 2 images in bmp format (.0004%).

RESULTS

Generalized model of CCI

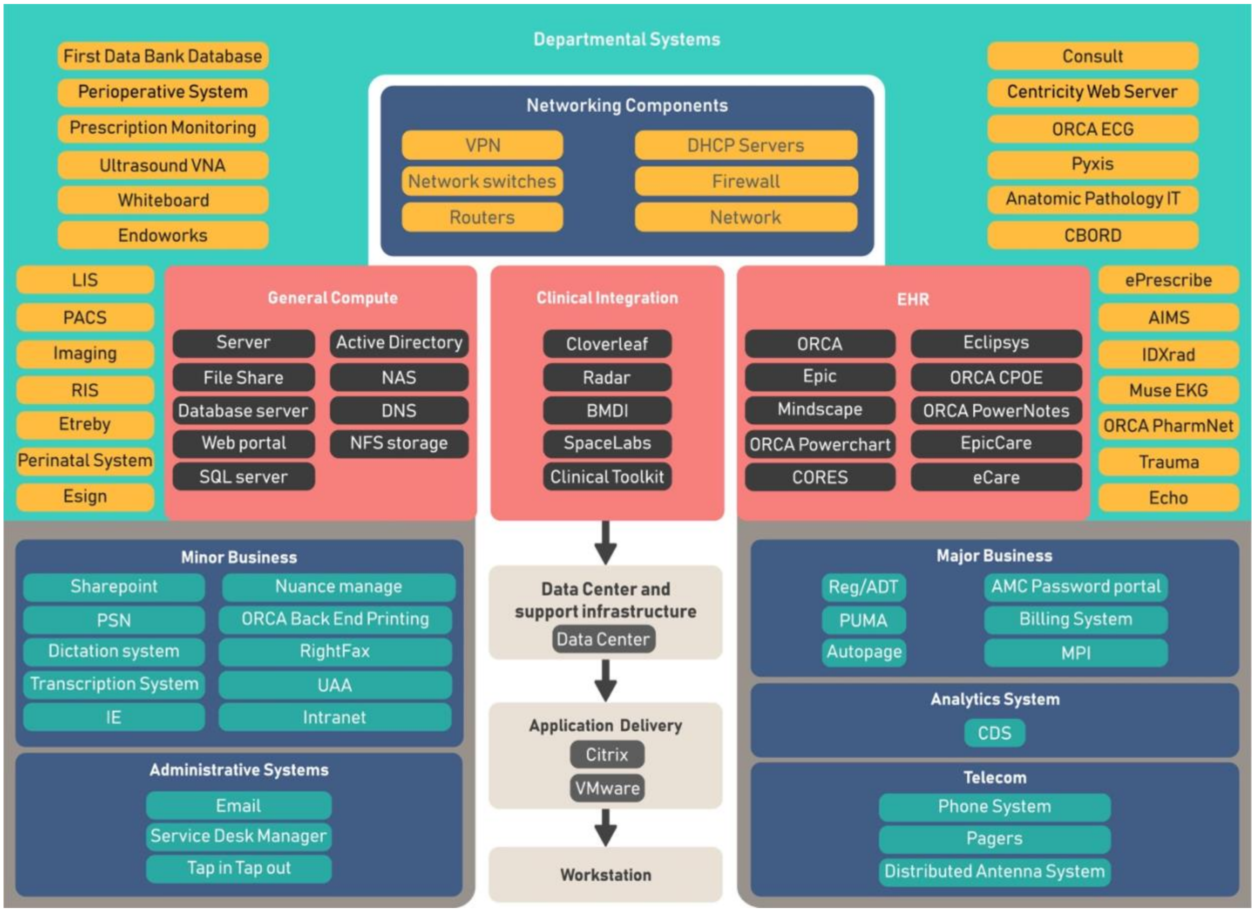

Legend : LIS – Laboratory Information System, PACS – Picture Archiving and Communication System, RIS – Radiology Information System, PSN – Patient Safety Net, IE – Internet Explorer, UAA – User Accounts and Authentication, VNA – Vendor Neutral Archive, VPN – Virtual Private Network, DHCP – Dynamic Host Configuration Protocol, BMDI – Bedside Monitor Device Interface, ORCA – Online Record of Clinical Activity, CORES - Computerized Rounding and Sign out, CPOE – Computerized Provider Order Entry System, ECG – Electro Cardiogram, AIMS – Anesthesia Information Management System, ADT – Admit Discharge Transfer, CDS – Clinical Decision Support System

Results from the analysis

A component as we define it here can either be a software application or hardware device that performs a specific function in the CCI. Since the dataset was large with more than 200 individual components, results were aggregated into 13 groups based on the common function served by the group, as assessed by domain experts within UW Medicine. Planned downtimes are recorded as scheduled. Outages are recorded as unscheduled. User impact is an estimate and is reported for available data. A total of 2725 unique records from the pool of 5019 files were isolated. 1614 incidents were scheduled (59.2%) amounting to a cumulative scheduled downtime of 3673.4 hours and 1111 incidents were unscheduled (40.8%) for a cumulative unscheduled downtime of 2972.6 hours.

All events

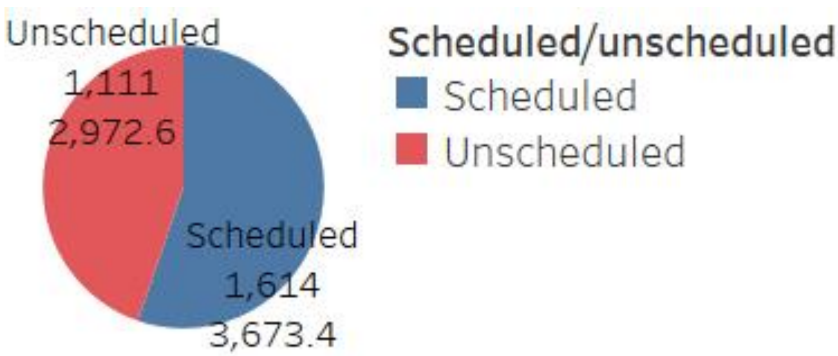

A yearly timeline for downtime periods is also obtained.

Timeline of events

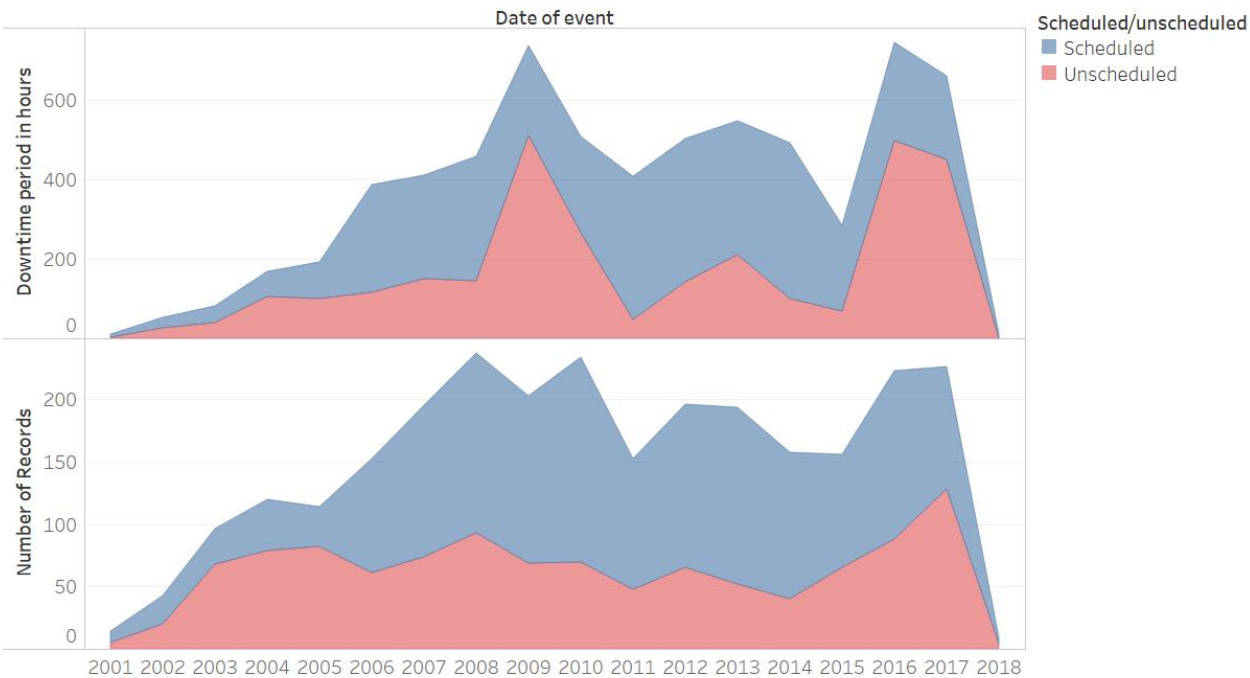

By absolute count (number of records) the top 5 components that failed most often are: Network (393 incidents, 59.5% unscheduled) ORCA (inpatient EHR) (372 incidents, 49.5% unscheduled) Epic (228 incidents, 12.3% unscheduled), Mindscape (130 incidents, 60 % unscheduled) and Data Center (123 incidents, 37.4% unscheduled)

Top 5 by Absolute Count

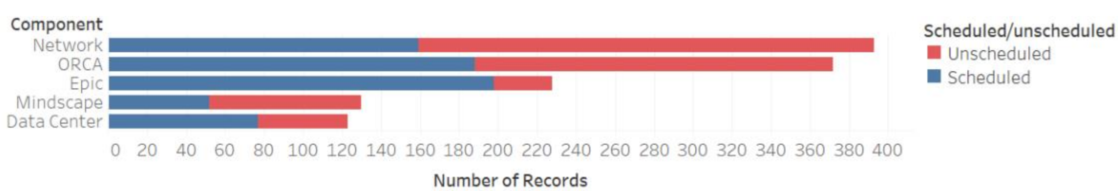

By cumulative amount of downtime spent in hours, the top 5 components are: network (1048.7 hours, 70.7% unscheduled), ORCA (977.3 hours, 54.5 % unscheduled), Epic

(567.8 hours, 5.3 % unscheduled), Cloverleaf (interface engine) (417.5 hours, 87.6 % unscheduled), Imaging (257.6 hours, 94.3 % unscheduled)

Top 5 by Downtime

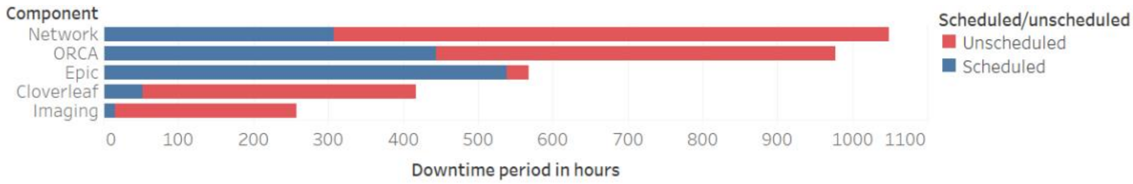

Examining user impact, the most impactful components according to the cumulative downtime spent are: ORCA (284.8 hours, under 5 users), Cloverleaf (263.5 hours, under 200 users), Imaging (205.8 hours, under 50 users), and Network (193.9 hours under 50, and 193.4 hours under 10 users)

User Impact Matrix

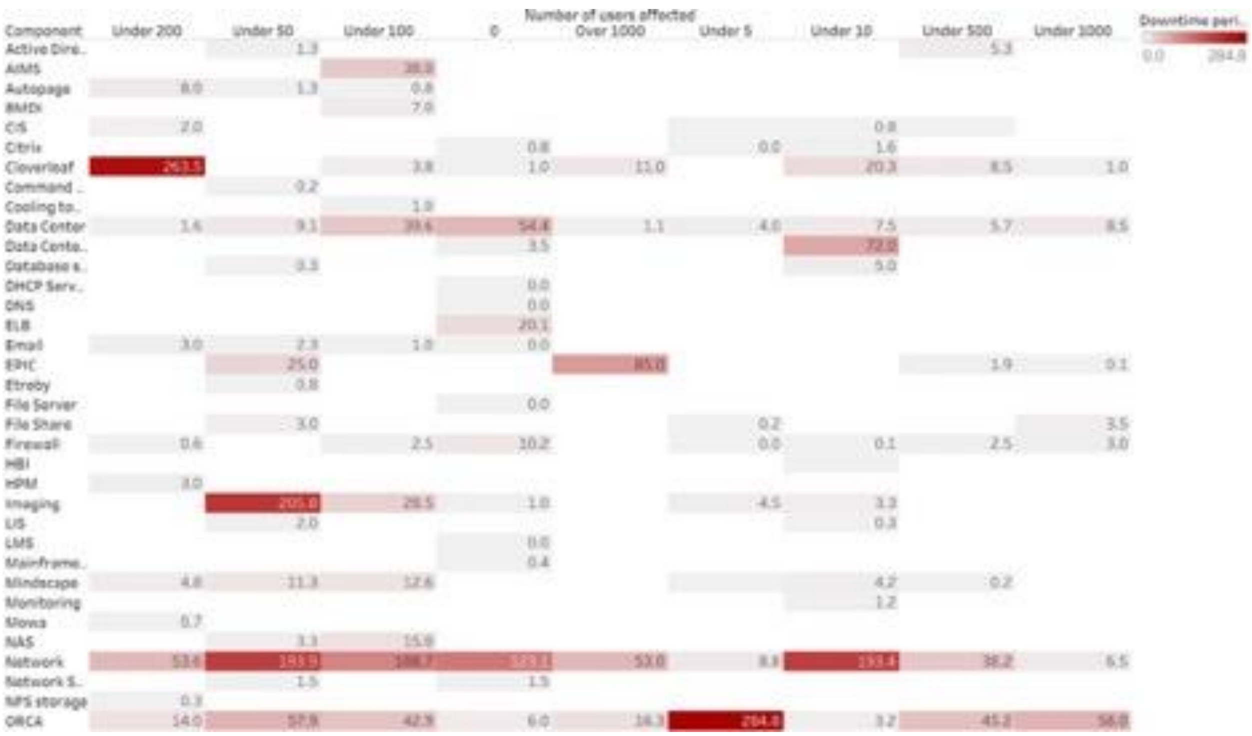

For the data with user impact estimates, cumulative EHR downtime (687.3 hours) is less than cumulative downtime for small impact non EHR CCI failures (1131.1 hours).

EHR components

Components included in this group support core EHR functions. These include CPOE (Computerized Practitioner Order Entry), results review, clinical documentation applications, patient portals, web-based systems to manage patient problem lists, diagnoses, electronic medication administration record, emergency room order entry and physician documentation, sign offs, rounding, and other workflows.

COUNT

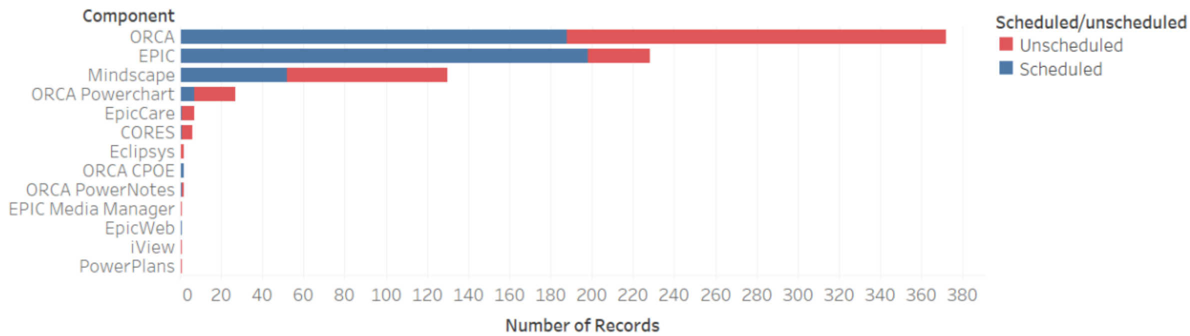

PERIOD

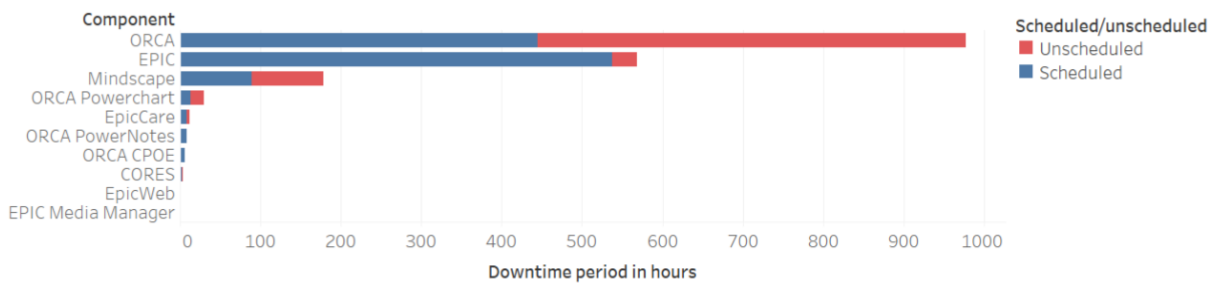

TIMELINE

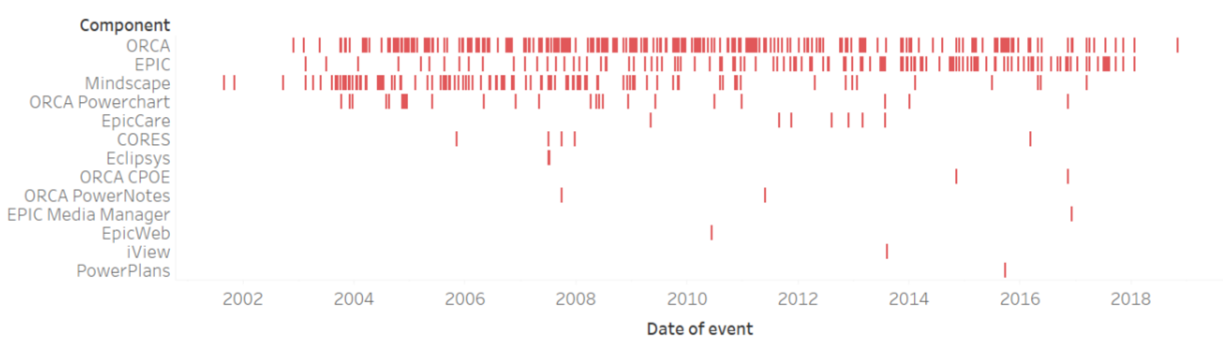

USER IMPACT

	Number of users affected									Downtime peri..
Component	Under 200	Under 50	Under 100	0	Over 1000	Under 5	Under 10	Under 500	Under 1000	0.1 – 284.8
EPIC		25.0			85.0			1.9	0.1	
Mindscape	4.8	11.3	12.6				4.2	0.2		
ORCA	14.0	57.9	42.9	6.0	16.3	284.8	3.2	45.2	58.0	
ORCA Powerchart		2.3	10.7				1.3			

Legend : CPOE – Computerized Provider Order Entry System, ORCA – Online Record of Clinical Activity

ORCA (Online Record of Clinical Activity) is the term used in UW Medicine for an electronic medical system based on the application suite from Cerner (Kansas City, MO). It is a suite of applications that include the main ORCA application, ORCA PowerChart, ORCA CPOE, ORCA PowerNotes, and PowerPlans and also other applications developed internally and embedded within or linked to PowerChart. Since 2003, UW Medicine has been using Cerner products for core inpatient workflow needs. The main ORCA Millennium application (represented as ORCA) reported about 372 instances of downtime, 184 were unscheduled (49.5%). Of the 977.3 hours of cumulative downtime, 444.4 hours (54.5%) was unscheduled. Most ORCA downtime was seen in under 5 users with a cumulative downtime of 284.8 hours in this user group. This was largely because of three events. ORCA PowerChart is an application that provides referring providers with read only access to EHR via VPN access. The information is real time encounter information and patient data. There were 27 instances of PowerChart failure, 20 were unscheduled (74.1%). The cumulative downtime was 29.8 hours of which 16.3 hours were unscheduled (54.7%). ORCA CPOE enables electronic entry and communication of orders, usually fulfilled by departmental systems or by nurses, therapists, and other health care team members. ORCA CPOE had 2 instances of failure, both of which were scheduled with a cumulative downtime of 5.5

hours. ORCA PowerNote is a note creation tool used to create clinical documentation. It had 2 instances of failure, both scheduled, with a cumulative downtime of 8 hours. PowerPlans are collections of order sets ; this failed 1 time, and was unscheduled. Our clinical image viewer iView (Cerner) accounted for 1 incident of failure. CORES (Computerized Rounding and Sign out) is a locally developed web-based system used to manage patient lists and sign out, store entered diagnoses, problem lists and medications[23]. There were 6 incidents with this component, 5 of which were unscheduled (83.3%) with a cumulative downtime of 2.3 hours. Before UW transitioned to Cerner in 2003 MINDscape (a locally-developed web results review application), was used for results review. MINDscape accounted for 130 events[24] of which 78 were unscheduled (60%). From a cumulative downtime of 177.9 hours, 89.2 hours were unscheduled (50.2%). Since 1997, UW Medicine has been using EpicCare (Epic Systems, Verona, WI) as the ambulatory care EHR in a progressively larger number of clinics. This has been expanded to outpatient specialty care in May 2014. EpicCare includes outpatient CPOE, clinical workflow and messaging to support a patient portal. There were about 228 incidents of Epic EHR failure of which only 30 were unscheduled (13.2%). From a cumulative downtime of 567.8 hours, only 30 hours were unscheduled (5.3%). There were 7 instances of EpicCare failure, 6 of which were unscheduled (85.7%), and 1 incident of EpicWeb downtime which was scheduled. eCare, the patient portal service, is a website for patients to access their diagnoses, lab results and communicate with their physicians. There were 4 incidents of downtime with eCare Epic Media Manager is used to upload patient related notes that providers may have composed offline or lab results or other clinical media. There was 1 incident of EPIC media manager downtime, which was unscheduled for 0.3 hours.

Departmental Systems

Components that operationalize care delivery processes are included in this category. Considering the broad scope of clinical services that constitute patient care at our institution, it is not surprising that this category has the most number of components, 26, in all.

COUNT

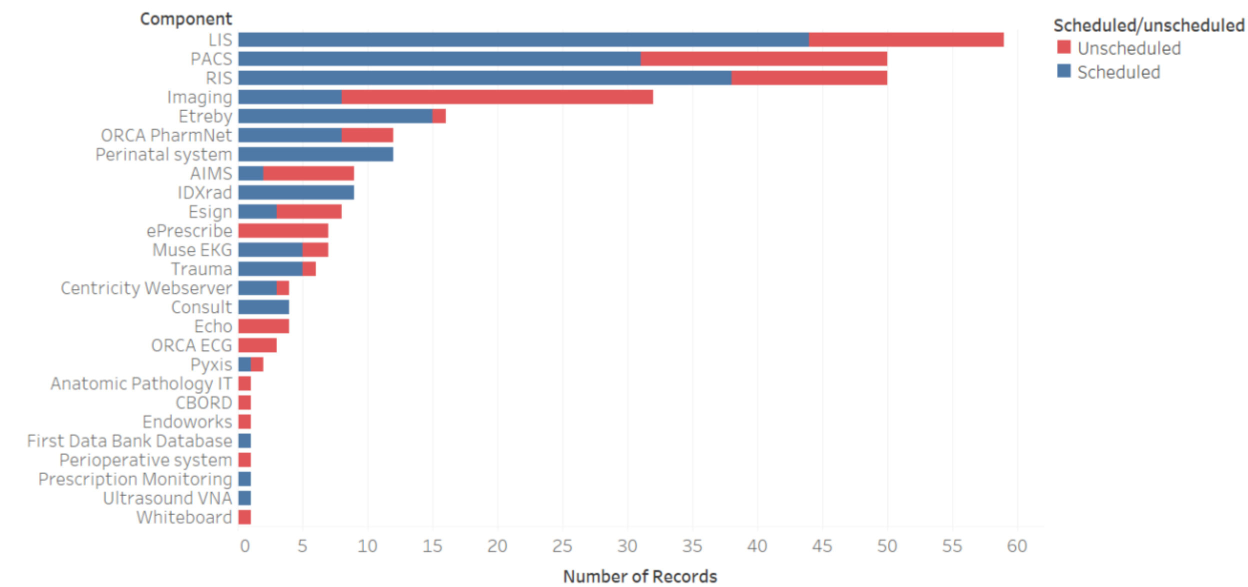

PERIOD

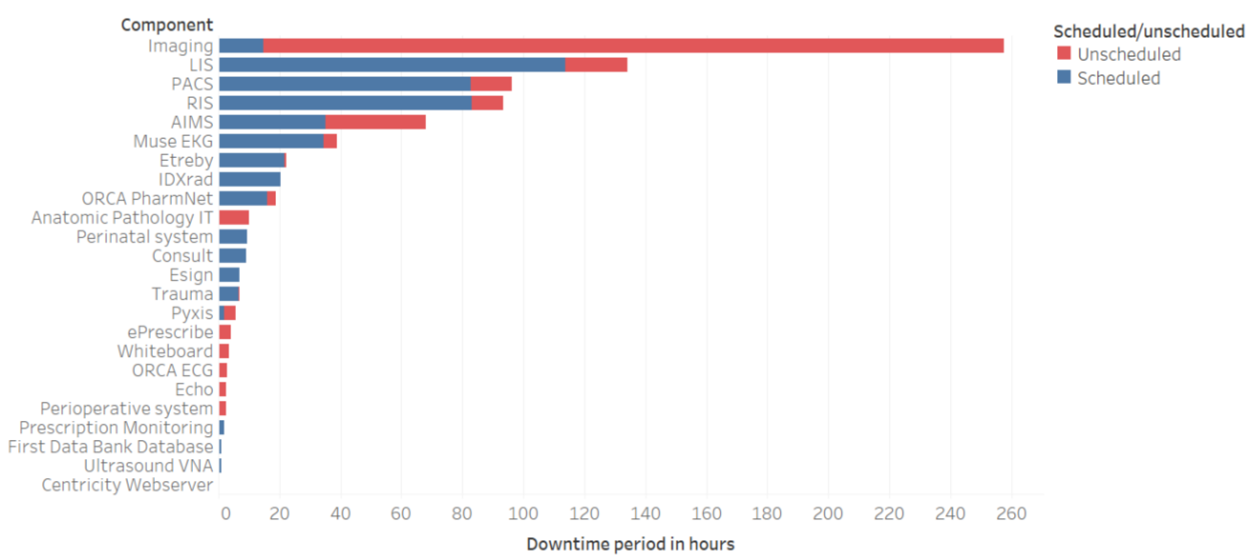

TIMELINE

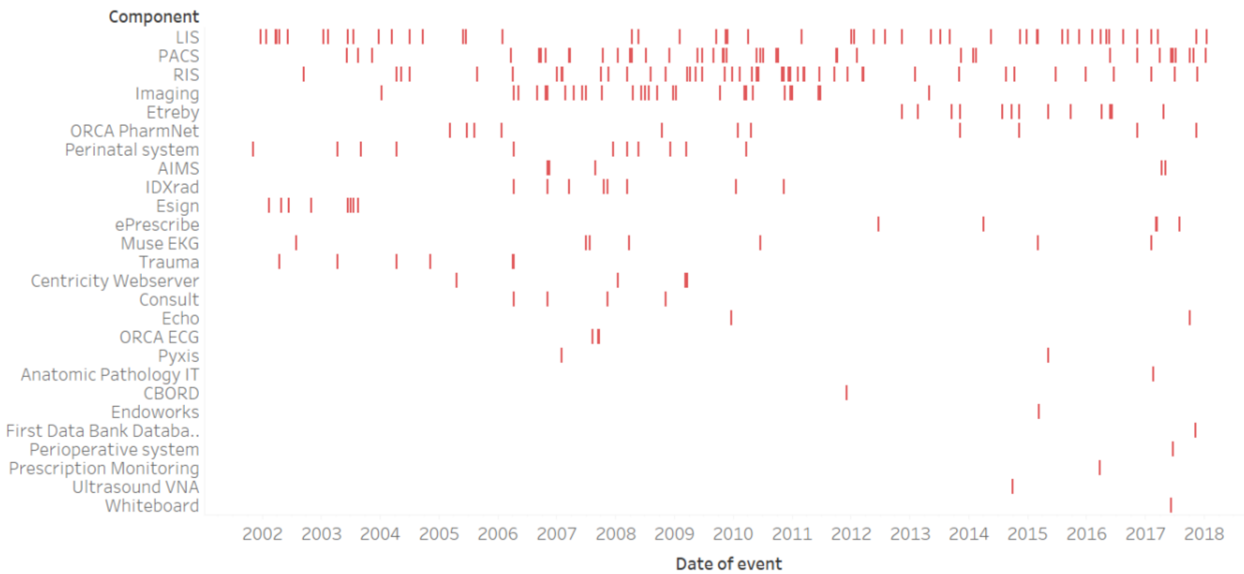

USER IMPACT

	Number of users affected					
Component	Under 50	Under 100	0	Under 5	Under 10	Under 500
AIMS		38.0				
Etreby	0.8					
Imaging	205.8	28.5	1.0	4.5	3.3	
LIS	2.0				0.3	
ORCA ECG	2.0	0.8				
ORCA PharmNet	1.5	1.3				
PACS	0.5	0.5				
Perinatal system		0.2				
Pyxis		3.5				
RIS						7.0
Trauma				0.2		

Downtime peri..
0.2 205.8

Legend : Anesthesia Information Management System, LIS – Laboratory Information System, ECG – Electro Cardiogram, PACS – Picture Archiving and Communication System, RIS – Radiology Information System,

The LIS (Laboratory Information System) record, manages and stores laboratory data by sending orders to lab instruments, tracking lab orders, and recording results in a searchable database. [25] Our LIS is licensed from Sunquest Inc and had 59 reported instances of downtime. Of these, 44 were scheduled (74.6%). The total cumulative downtime was 134.3 hours of which, 113.8 hours were scheduled 84.7%. Anatomic

Pathology is a component that streamlines pathology sample processing. We have 1 instance of downtime in our records. It was an unscheduled event for 10 hours. PACS (Picture archiving and communication system) is a technology that allows storage and access to clinical images such as ultrasound, magnetic resonance images, nuclear medicine imaging, computerized tomography scans, endoscopy scans, mammograms, digital radiographs and histopathology scans among other imaging modalities.[26] Our PACS licensed from General Electric Company (GE), had 50 incidents of which 31 were scheduled (62%). The cumulative downtime was 96.1 hours of which 82.6 hours were scheduled (86%). Our imaging service, which performs a similar function as PACS, had about 32 incidents of which 24 incidents were unscheduled (75%). From a total of 257.6 hours of cumulative downtime, 242.9 hours were unscheduled (94.3%). Most user impact was seen in the under 50 user range (cumulative downtime of 205.8 hours). The RIS (Radiological information system) complements PACS. It facilitates the electronic management of imaging departments in the radiology practice. Our RIS licensed from GE had about 50 instances of downtime, 38 of these were scheduled (76%). From a cumulative downtime of 93.5 hours, 83 hours were scheduled (88.8%). Our older RIS was IDXrad (which was acquired by GE). It had about 9 instances of downtime all of which were scheduled for a cumulative downtime of 20.3 hours. Centricity web server (GE co) facilitates the web interface for clinicians to review clinical images. There were 4 instances of Centricity web server downtime of which 3 were scheduled events (75%). Etreby is a retail pharmacy system to help manage outpatient pharmacy and monitor patients post discharge. It is now a Cerner product. There were 16 instances of Etreby Retail pharmacy downtime in our records, of which only 1 was unscheduled (6.3%). From a 22.2-hour cumulative downtime, the unscheduled event lasted 0.8 hours.

The retail pharmacy system we used before Etreby was ORCA PharmNet which had 12 incidents of downtime, 8 of which were scheduled (66.7%). 18.7 hours of cumulative downtime was accrued, of which 16 hours were planned (85.3%). Perinatal system is a software (licensed from GE Centricity Perinatal System) that is designed specifically for the perinatal suite to monitor labor, maternal and child health parameters. It is also used in the NICU setting. There were about 12 instances of downtime, all of which were scheduled for a cumulative downtime of 9.3 hours. Esign is a locally developed software that allows providers to digitally sign orders. Our dataset contained about 8 incidents related to eSign downtime, 5 of which were unplanned (62.5%). ePrescribe (eRx) is a system developed by Allscripts Inc that allows a clinician to generate, transmit and fill a medical prescription. Prescriptions can also be renewed electronically. There were about 7 eRx related incidents, all of which were unscheduled for a cumulative downtime of 4.1 hours. AIMS stands for Anesthesia information management system, which allows the collection, storage and relay of patient information during the perioperative period[27]. There were 9 instances of AIMS downtimes in our records, 7 of which were scheduled (77.8%). From a cumulative downtime of 68 hours, 35 hours were scheduled (51.5 %)

Muse EKG is a cardiology information management system licensed from GE. It relays EKG information securely to providers and establishes connectivity with Vendor Neutral Archives (VNA) to store data. We have 7 incidents of Muse EKG downtime in our records, 5 of these were scheduled (71.4%). A cumulative downtime of 39 hours was seen of which 34.5 hours was scheduled (88.5 %). Included in the Centricity suite of applications is Echo, a component specialized for visualizing echocardiography data.

There were about 4 instances of Echo downtime in our records, all of which were unscheduled for a total of 2.5 hours. Before we switched to Muse EKG we used, ORCA ECG application, which had 3 incidents of failure, all of which were scheduled for a total of 2.8 hours. Pyxis is an automated medication dispensing system licensed from Becton, Dickinson and Company. It supports decentralized medication management and pharmacy workflows. There were 2 instances of Pyxis failure in our records, 1 unplanned incident of 3.5 hours and 1 planned incident of 2 hours. Consult was a locally developed application to manage referrals and provider consults. There were 4 incidents, all of which were scheduled for a total of 9 hours. CBORD (Roper Technologies Inc.) is a food services application used to plan patient meal plans. We recorded 1 instance of downtime, which was unscheduled. First data bank is a drug database application to reference drug products, dispense medications, manage formulary, perform drug pricing analysis and process medical insurance claims to support eRx. We have 1 instance of failure in our records, which was scheduled maintenance. The perioperative system used at our institution is called Horizon Surgical Manager (McKesson Information Solutions) helps schedule patients for surgery, monitor their clinical data and discharge them when suitable. There were 1 instance of downtime in our database. It was an unscheduled event for 2.5 hours.

A Vendor Neutral Archive (VNA) stores clinical documents and images in a standardized format for intersystem accessibility. Our ultrasound VNA (licensed from GE) had had 1 instance of downtime which was scheduled for 1 hour. Electronic whiteboards help manage patients by tracking them, monitoring their flow and facilitating provider to provider communication. There was 1 instance of electronic whiteboard downtime at our institution. It was unscheduled and for 3.5 hours.

Endoworks was an Endoscopy information management system developed by Olympus America Inc. There was 1 instance of downtime in our records, which was unscheduled.

Clinical Integration

The components included in this category integrate data flows to and from the EHR, Reg/ADT, Imaging and other departmental systems. Without these integration systems, information exchange would be point-to-point, demanding a lot of infrastructure hardware. Interface engines make one to many connections possible, so clinical data integration is easier.

COUNT

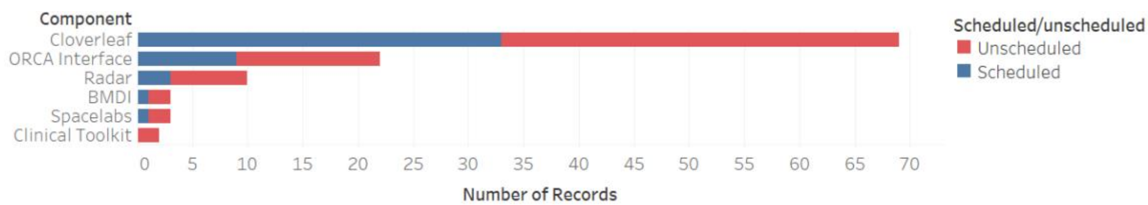

PERIOD

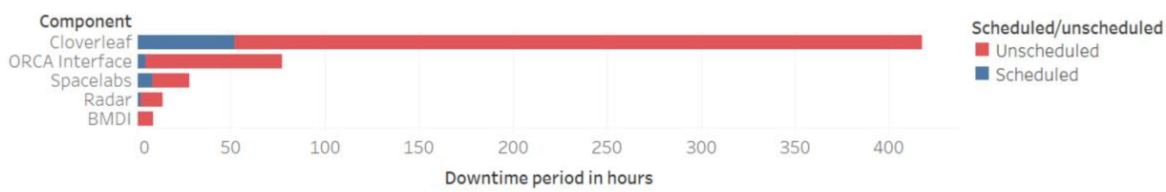

TIMELINE

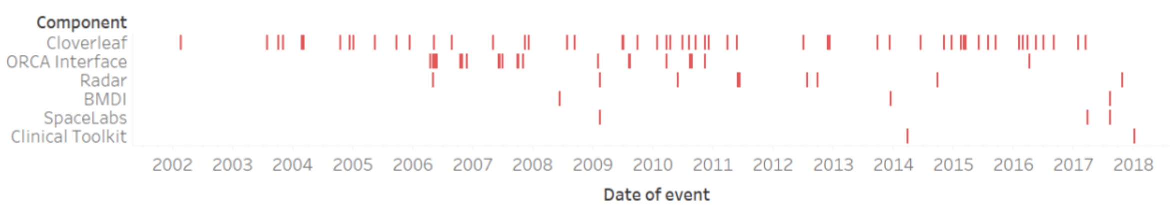

USER IMPACT

	Number of users affected								
Component	Under 200	Under 50	Under 100	0	Over 1000	Under 5	Under 10	Under 500	Under 1000
BMDI			7.0						
Cloverleaf	263.5		3.8	1.0	11.0		20.3	8.5	1.0
ORCA Inter..	2.5	25.8	7.0		2.5	1.5	13.2		
Radar						0.7	6.3		

Downtime peri.. 0.7 263.5

Cloverleaf (Corepoint Health Inc) is one of the interface engines currently used at our institution. It ensures the exchange of clinical data between multiple components using messages following HL7 standard and others. We have had 69 instances of interface engine failure of which 36 events were unplanned (52.2%). The cumulative downtime was 417.5 hours of which 365.8 hours were unscheduled (87.6%). The user impact of Cloverleaf failure was seen in the under 200 user range for a cumulative downtime of 263.5 hours. ORCA interface, and Radar were our interface engines used before we adopted Cloverleaf technology. There were seen about 22 instances of ORCA interface failure, 13 of which were unscheduled (59.1%). The cumulative downtime was 77.4 hours of which 73 hours were unscheduled (94.3%). We noticed 10 instances of Radar failure, 7 were unscheduled (70%). A cumulative downtime of 13.9 hours was seen of which 11.9 hours (85.6%) was unscheduled. BMDI stands for Bedside Monitoring Device Interface. In the ICU settings, physiological parameters, wave forms, device settings and alarm data are directly imported into the EHR with this technology. Our data has 3 instances of BMDI downtime, 2 of which were unscheduled (66.7%) for a total downtime of 8.5 hours. SpaceLabs was our old BMDI which had 3 instances of downtime, 2 of which were unscheduled (66.7%) for a total downtime of 28 hours. Clinical toolkit is a web page that contains links to various other resources in the organization. There were about 2 instances of toolkit downtime, both of which were unscheduled.

Business Major

The components included in this category have a major impact on business operations in addition to clinical impact. These systems include patient registration systems, password portals, billing systems and the master patient index among others.

COUNT

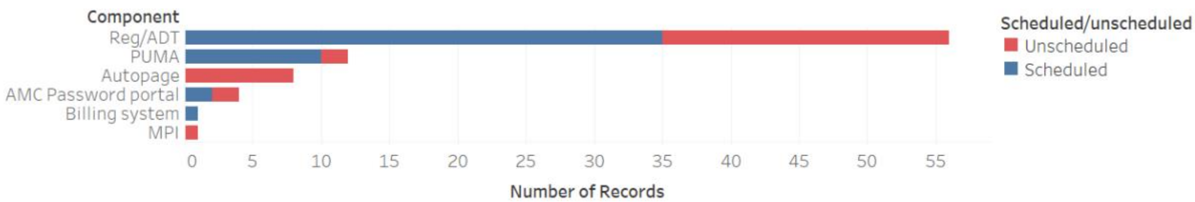

PERIOD

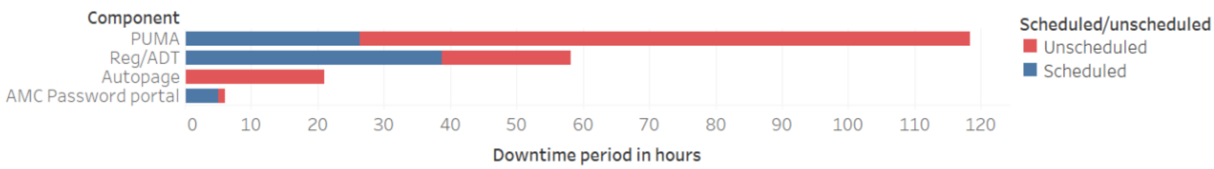

TIMELINE

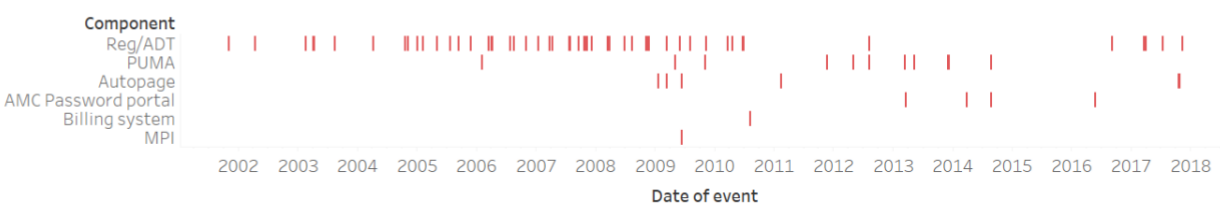

USER IMPACT

	Number of users affected					
Component	Under 200	Under 50	Under 100	Under 10	Under 500	Under 1000
Autopage	8.000	1.250	0.750			
PUMA				3.000		
Reg/ADT	9.500			1.250	5.750	2.500

Downtime peri..
0.750 9.500

Legend : ADT – Admit Discharge Transfer, MPI – Master Patient Index, PUMA – Physician Utilization Management Advisor

Reg/ADT represents patient registration and ADT module (Admit-Discharge-Transfer). It carries demographic information in addition to trigger event data (patient admission, discharge, transfer, registration etc.). Our old ADT system had 56 instances of downtime, 35 of which were scheduled (63%). A cumulative downtime of 58.2 hours was noted, of which 38.7 hours were scheduled (66.5%). Most user impact is in the under 200 users range. PUMA stands for Physician Utilization Management Advisor. It is a program from the Veterans Affairs to see if patients are meeting the standard of care. We have seen about 12 instances of failure of which 2 were unscheduled (16.7%) Autopage is a computing script that reduces the radiologist's effort to page clinical data to referring clinicians[28]. There were about 8 instances of Autopage downtime at our institution, all of which were unscheduled for a cumulative downtime of 21 hours. The password portal is a locally developed application that verifies provider credentials and providers a place for access management and service functions. There were 2 scheduled (5 hours total) and 2 unscheduled instances (1 hr total) of password portal failure downtime. One of the most business-critical functions is the billing system, which had 1 instance of downtime which was scheduled for 1 hour. MPI stands for Master Patient Index. It is an enterprise wide patient database containing unique identifiers, so each individual is represented only once across all of the organization's systems[29]. Failures of MPI are fortunately uncommon, represented by 1 instance in our records, which was unscheduled. Application delivery refers to those components that enable a personalized desktop for each clinician user at any workstation in the institution. This technology called virtual desktop infrastructure (VDI) helps run a simulated machine that lives on a server in the datacenter, allowing IT departments to reduce costs and centralize desktops. Citrix Systems Inc provides VDI at our institution. We have had 37 instances of VDI downtime. We also have 1 instance of VMware failure.

Networking

Hardware and software components that facilitate inter-component data transfer, resource sharing and communication through physical and wireless connections are included in this category. In a hospital setting the scale of operations can range from connecting basic peripherals and mobile devices at terminal workstations to running massive data centers.

COUNT

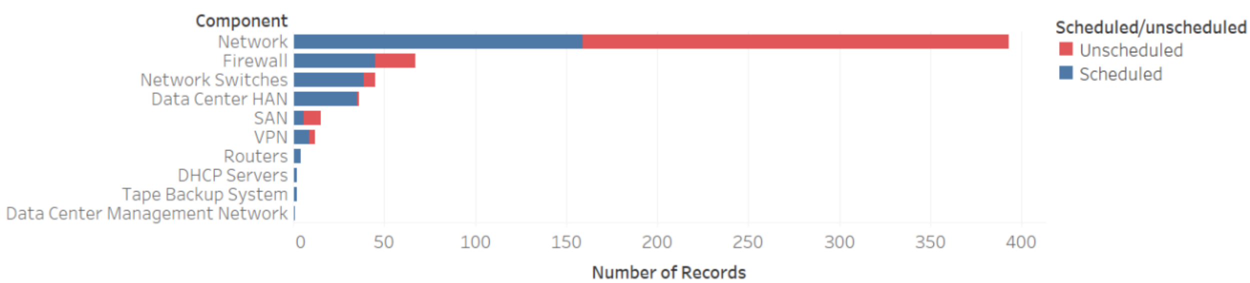

PERIOD

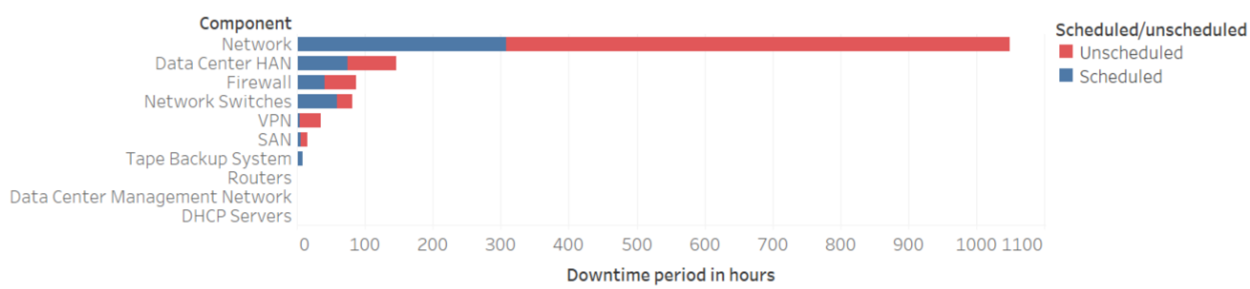

TIMELINE

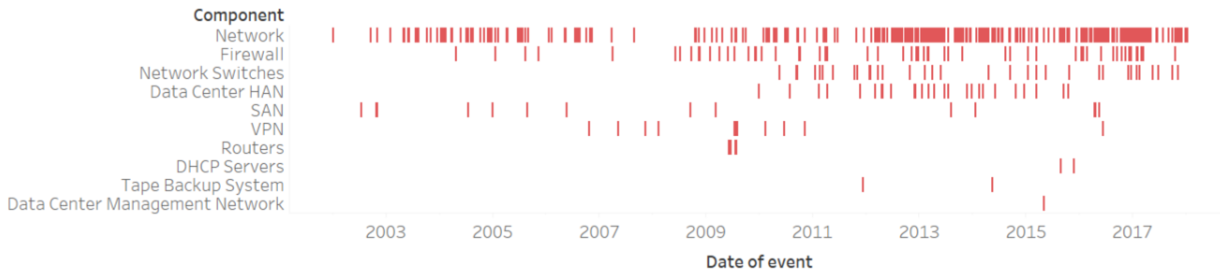

USER IMPACT

	Number of users affected								
Component	Under 10	0	Over 1000	Under 5	Under 50	Under 100	Under 200	Under 500	Under 1000
Data Center HAN	72.0	3.5							
DHCP Servers		0.0							
Firewall	0.1	10.2		0.0		2.5	0.6	2.5	3.0
Network	193.4	123.1	53.0	8.8	193.9	108.7	53.6	38.2	6.5
Network Switches		1.5			1.5				
SAN		0.5	4.0			3.5	0.2		
VPN					30.0	0.3	1.0		

Downtime peri..
0.0 193.9

Legend : HAN – High Availability Network, DHCP – Dynamic Host Configuration Protocol, SAN – Storage Area Network

Network includes primarily the wired and wireless connections that clinicians use to access the EHR, departmental systems and clinical toolkits. There were about 393 instances of network downtime representing the largest downtime category. Of these 234 incidents were unscheduled (59.5%). The cumulative downtime for network related outages was 1048.7 hours of which 741.4 hours were unscheduled (70.7%). Most user impact (up to 194 hours of cumulative downtime each) is noticed in under 10 users and under 50 users. However large outages are also seen: 38.2 hours of Under 500 users and 6.5 hours of under 1000 users. Firewalls allow network access requests from only safe sources, blocking unsafe or unknown sources. Our institution had 52 instances of firewall related downtime, 45 of these were planned outages (67.2 %). Of the 87 hours of cumulative firewall outage, 41.1 hours were planned (47.3 %). User impact was seen in Under 100 (2.5 hours), Under 200 (0.6 hours), Under 500 (2.5 hours) and under 1000 users (6.5 hours). Datacenter high availability network (HAN) represents networking components at a datacenter that ensure uptime and reliability. Uptime refers to the amount of time a datacenter is operational[30]. Our HAN had about 36 instances of downtime, 35 of them (97.2%) were planned. The cumulative outage was 146 hours of which 74 were scheduled (50.7%). There were one event of 72 hours unplanned downtime. User impact was low (under 10 users). Data Center Network Manager (DCNM, Cisco Systems Inc) is a network management platform to manage HAN. There were 1 instance of DCNM downtime. Storage area network is a type of local area network that enables bulk storage of data and handles large data transfers using high end servers, interconnect technology and multiple disk arrays. Our SAN had 15 instances of downtime, 9 of which were unscheduled (60%). The cumulative downtime was 15.5 hours, 10.4 hours were unscheduled (67.3%). SAN impacts end users in the 100-200 users range. A tape backup system[31] is a traditional sequential back up procedure used in enterprises because of its archival stability. There were 2 instances of tape backup system failure, both of which were scheduled for a cumulative maintenance downtime of 8 hours. VPN stands for virtual private network, a system that allows clinicians to login remotely. There were 12 instances of VPN downtime at our organization, 9 of these (75%) were scheduled. The cumulative downtime was 35.4 hours, of which 31.3 hours were unscheduled (88.4%). User impacts were mostly seen in the Under 50 range. A network switch is a networking hardware component that connects devices on a network by packet switching to determine the ultimate destination of the data packet. There were 45 instances of network switch failure. 39

incidents were planned outages (86.7%). The cumulative downtime was 82.3 hours, 59.5 hours were scheduled (72.3%). User impact is under 50 users range. A router is a networking hardware component that forwards data packets between computer networks. It works like a post box sorting the packets ultimate destination depending on its IP address. There were 4 router related downtimes, all of which were planned maintenance downtimes for a cumulative 1.9 hours. Dynamic Host Configuration Protocol servers dynamically assign unique network configurations (IP address) to each device on the network so they can communicate with each other and other IP networks. There were 2 instances of DHCP server downtime both of which were scheduled.

Data center and support infrastructure

A data center is a facility within an enterprise that houses and maintains backend IT infrastructure such as mainframes, servers and databases to manage store and disseminate data across the enterprise[32]. Support infrastructure refers to components that contribute to sustain a high level of availability, reliability and security to the data center. These typically include a UPS (uninterruptible power supply), environmental controls, and physical security systems) Our database contained about 123 incidents of Data Center downtime, 77 of which were scheduled maintenance events (62.6%). Within a cumulative data center downtime of 209.8 hr, 77.2% (162 hr) were scheduled and because of the inherent redundancies in a data center operation, most user impact was negligible as noted in the user impact table.

COUNT

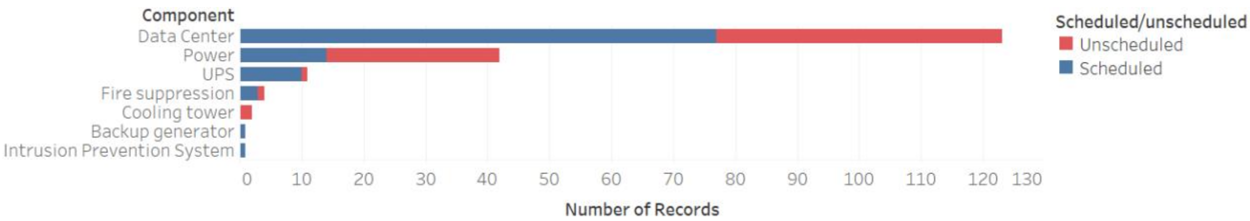

PERIOD

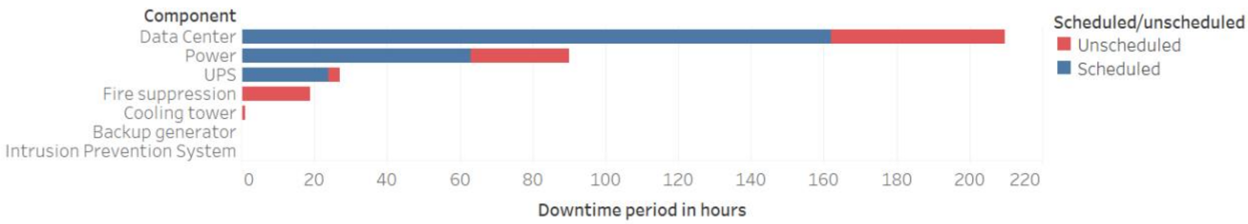

TIMELINE

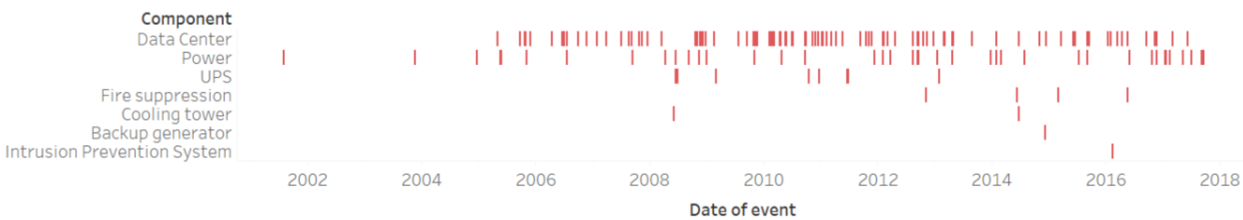

USER IMPACT

	Number of users affected								
Component	Under 200	Under 50	Under 100	0	Over 1000	Under 5	Under 10	Under 500	Under 1000
Cooling tower			1.00						
Data Center	1.60	9.05	39.55	54.40	1.05	4.00	7.50	5.68	8.50
Power				3.50					
UPS				10.00					

Downtime peri..
1.00 54.40

Legend: UPS – Uninterruptible Power Supply

Our power infrastructure includes power (42 downtime incidents) and 67% of these incidents were unscheduled. From a total of 90 hours of recorded power downtime, only 30% of it was unscheduled. UPS had about 11 downtime incidents, of which 10 were scheduled. It had 27 hours of cumulative downtime, 88.9% of which was scheduled. The backup generator had 1 downtime incident in 2014. Our environmental controls include fire suppression (4 downtime incidents, 1 of which was unscheduled (25%)).The cooling tower had 2 downtimes, both unscheduled) Our security infrastructure, the intrusion prevention system had 1 scheduled downtime event.

General compute

In this category we have components that facilitate general purpose computing, which is computationally intensive but serves the general computing requirements like a database server rather than perform a specific function like an EHR application.

COUNT

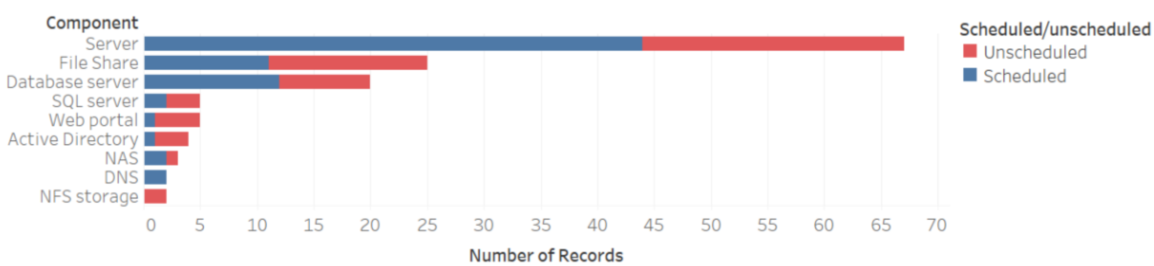

PERIOD

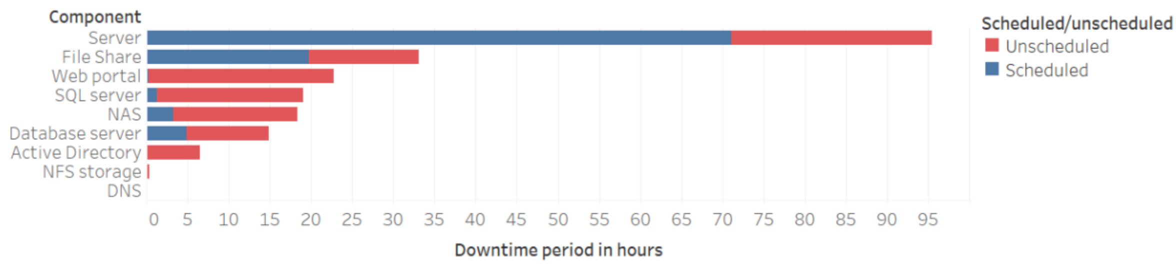

TIMELINE

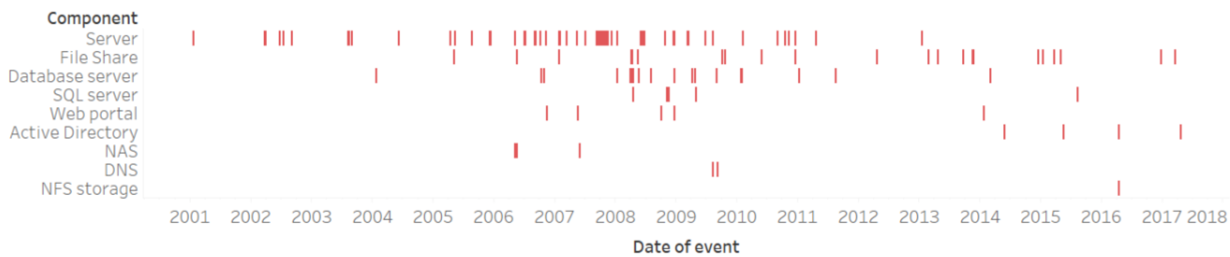

USER IMPACT

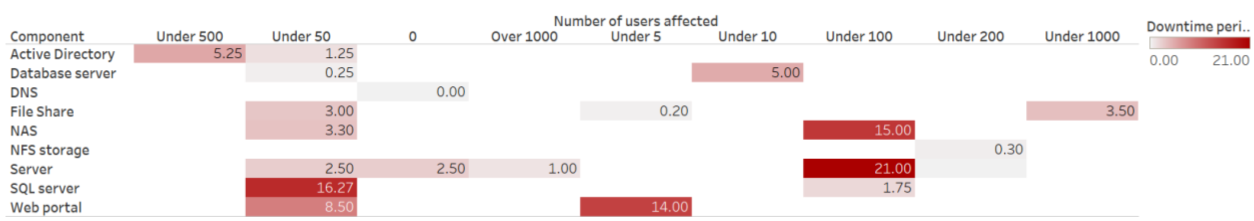

Legend – DNS – Domain Name Server, NAS – Network Attached Storage, SQL – Structured Query Language, NFS – Networked File Server

Servers are devices or programs that provide functionality (or services) to other specialized programs or devices called clients[33]. We use a typical client-server model to distribute computational resources and data. There are many types of servers:

computing servers, database servers, file servers, and application servers. Computing severs typically share computing resources such as processing power and random-access memory. In our database, we recorded 67 instances of computing server downtimes, of which 44 were scheduled (65.7%). The cumulative downtime recorded was 95.45 hours of which 71 hours were scheduled (74.4%). File servers share storage space over a network. There were 25 instances of file share server downtimes of which There were an almost even split of 14 unscheduled (56%) events. The cumulative downtime was 33.12 hours of which 13.37 hours or 40.3% was unscheduled.
There were 3 instances of NAS (network attached storage) downtimes, 1 of which was unscheduled with a cumulative downtime of 18.3 hours. The unscheduled incident lasted 15 hours. Database servers maintain and share organized tables of databases. There were 20 instances of database server downtimes, 8 of which were unscheduled. The cumulative downtime was 14.85 hours, 10 hours were unscheduled. There were 5 instances of SQL server downtimes, 3 of which were unscheduled (60%). The cumulative downtime was 19.02 hours of which, 17.77 hours were unscheduled (93.4%). Name servers are the nodes of the Domain Name System (DNS), which help in translating host names into an IP address There were 2 instances of DNS downtime, both of which were scheduled. Active directory is an example of a mail server directory service (Microsoft Inc) that is used to configure windows domain networks. There were 4 incidents in our database, 3 of which were unscheduled representing a cumulative downtime of 4.85 hours. Web servers host web pages called web portals. There were 5 incidents of web portal downtime in our database, 4 of which were unscheduled with a cumulative downtime of 22.5 hours.

Telecom

In this category components that are mainly used for telecommunications across our enterprise are included. Although there is significant overlap in modern communications technology with networked IT components We include phone lines, pagers and antenna systems in this category.

COUNT

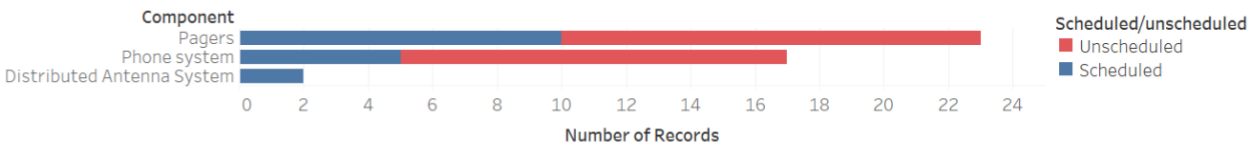

PERIOD

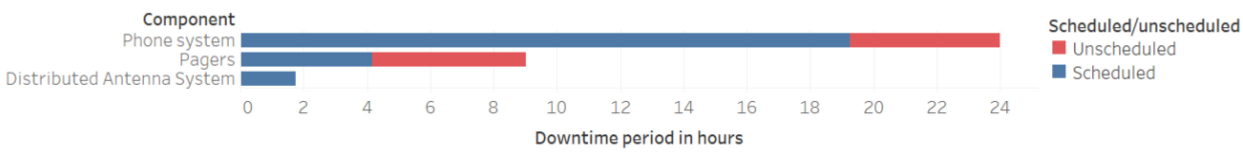

TIMELINE

Pagers are wireless telecommunications devices that receive text messages. They are still used today because some pagers rely on satellite communications making them more reliable and with improved reception in some areas in comparison with ground based cellular radio networks[35]. There were 23 incidents of pager downtime in our records, 13, 56.5 % of which were unscheduled. Of the 9.02 cumulative hours of pager downtime, 4.85 hours 53.8% was unscheduled. Phone system in our institution includes the telephone system, help desk phone system, teleservices, voice mail and SpectraLink systems. We have 17 incidents of phone line related downtime,12, 70.6 % of which were unscheduled. Our phone system had events in 2017 and 2018. It was down recently because of teleservices trunking problems. A Distributed Antenna System or DAS is a group of antenna nodes separated in space via a transport medium, connected to a common source that provides wireless services in an enterprise[35]. Our database has 2 recorded incidents of DAS downtime both of which were scheduled maintenances.

Business minor

In this category we include components that only have a minor impact on clinical operations at an institution. By minor we mean workarounds for failures of these components are easy and no mission critical functionality is lost.

COUNT

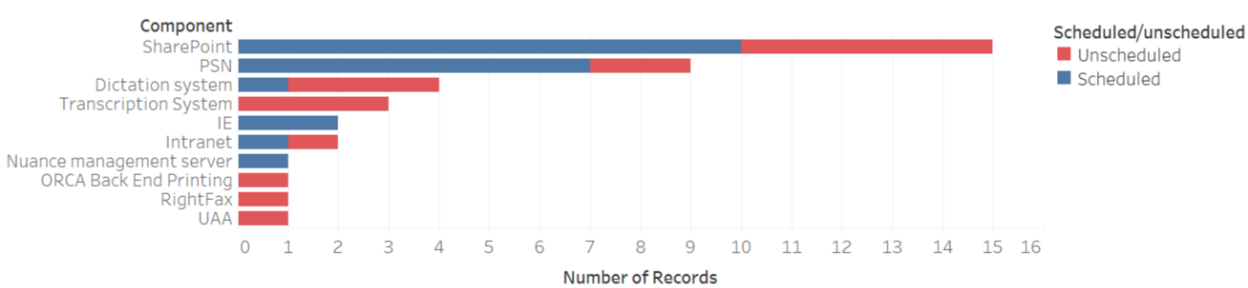

PERIOD

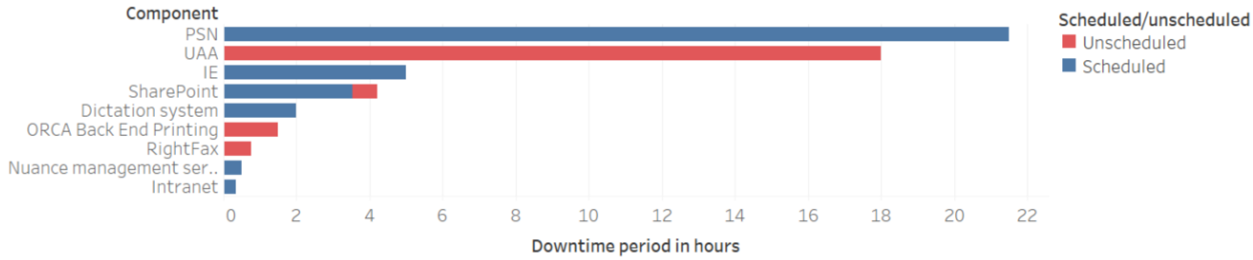

TIMELINE

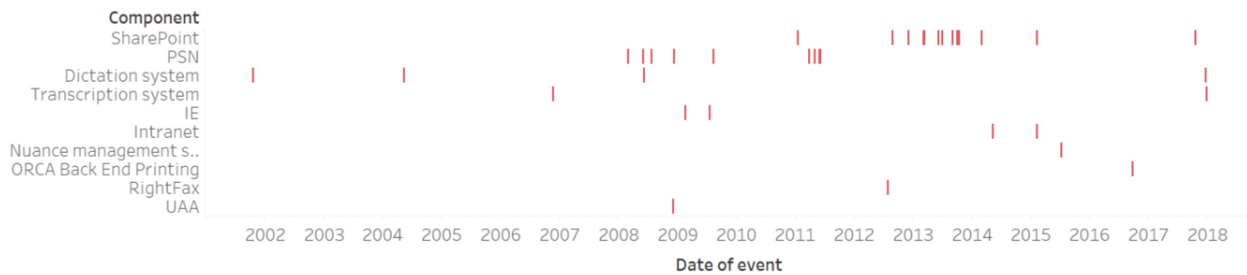

Legend : PSN – Patient Safety Net, IE – Internet Explorer, ORCA – Online Record of Clinical Activity, UAA – User Account and Authentication

SharePoint (Microsoft Corp) is a web-based document management and storage service that integrates with Microsoft office suite of applications. We noticed 15 instances of downtime in our database.10 of these were scheduled maintenance downtimes (66.7%). 4.21 hours of cumulative downtime was accrued, of which 3.54 hours (84.1 %) was scheduled. Patient Safety Net is a web based, real time safety intelligence software that is used for reporting patient, visitor, and staff safety (actual and potential) risks, errors, harmful behavior, and unsafe conditions. We recorded 9 incidents of PSN related downtimes. 7 of these incidents were scheduled 77.8% contributing to a 21.5-hour cumulative downtime. Dictation system and Nuance management server are voice recognition systems that recognize speech and help providers document clinical notes faster than typing. Our dictation system, Dragon medical client (Nuance Inc) had 4 instances of downtime, 3 of these were unscheduled

(75%). Since we have multiple clients running throughout our organization, we manage them from a single server called the Nuance management server. This server has had one scheduled downtime for 0.5 hours, according to our database. Transcription system is an interface for human driven transcriptions (medical transcription), usually outsourced. These were popular before machine driven transcriptions reached an acceptable level of accuracy. There were 3 instances of downtime all of which were unscheduled. Internet Explorer (Microsoft Corp) is the terminal desktop browser that allows end users to access web-based applications. Our records contain 2 instances of Internet explorer downtime, both of which were scheduled for a total period of 5 hours. Intranet is a local private network accessible by authorized staff. These networks are usually established within the Local Area Network or Wide Area Network. Intranets deliver shared content and resources. We recorded about 2 incidents of intranet downtime one of which was scheduled (50%) for 0.33 hours. ORCA backend printing and RightFax are centralized secure backend printing and fax services across the entire organization. These integrate with email, desktop and document management applications to enable high volume automated functions. There was 1 incident each of ORCA backend printing (Cerner Inc) downtime for 1.5 hours and RightFax (OpenText Corp) downtime for 0.75 hours, each. User Account and Authentication (UAA) server is an identity management service that authenticates users, issues tokens for client applications and helps manage user accounts, including setting passwords. Our database contained 1 instance of UAA server failure, causing provider lock out from the system. This was unscheduled and for 18 hours in duration.

Devices

In this category we include specific hardware components that enable human interfacing with data.

COUNT

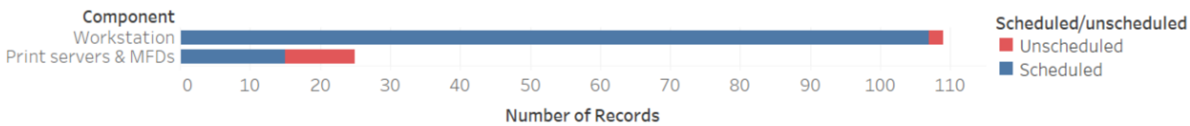

TIMELINE

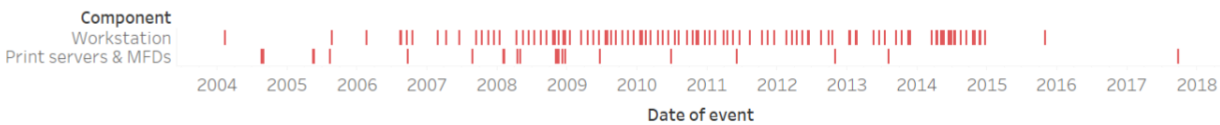

Legend : MFD – Multifunctional Device

Workstations are devices allow end users to access an EHR. These are desktop computers, usually running an enterprise version of Windows with applications installed to perform clinical and administrative duties. As such most (98.2%) downtimes were maintenance (Windows/Java updates) related. We recorded 110 instances of downtime. Print servers and MFDs (multi-function devices) can fail in addition to the backend printing applications we discussed earlier. We recorded 25 incidents in our database. The cumulative downtime was 29.20 hours of which 20.20 hours were unscheduled (69.2%).

Administrative systems

Components that are important to clinical and business operations at our institution are included in this category.

COUNT

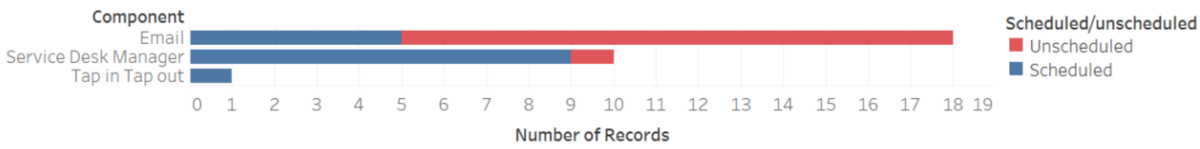

PERIOD

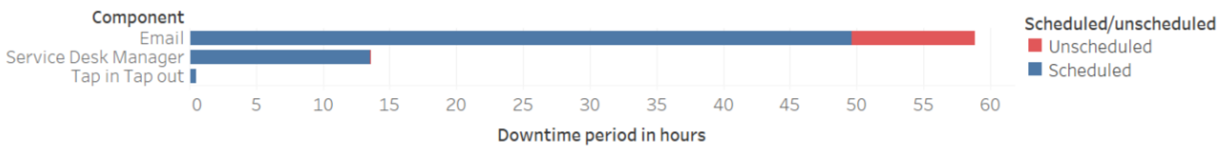

TIME LINE

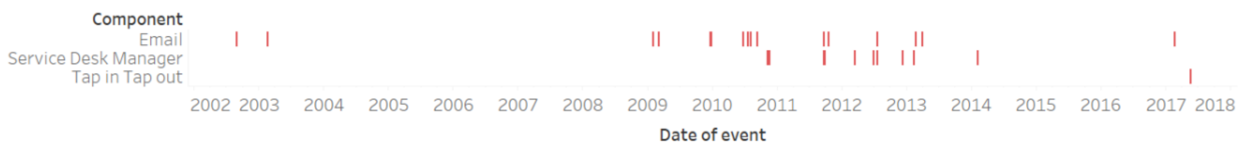

Email failures may manifest because of different component failures such as exchange server failure, SMTP failure (Simple mail transfer protocol) or network failure. Our database contained 18 instances of email downtime, 13 of which were unscheduled (72.2%). The cumulative email downtime was 58.85 hours of which only 9.25 hours or 15.7% was unscheduled. Service desk manager is a help desk ticketing system for providers and other end users to register their complaints with the IT system. There were 10 incidents of service desk manager downtime, only 1 event was unscheduled. A cumulative downtime of 13.52 hours was observed with the unscheduled event being only 0.02 hrs. in duration (1%). Tap in tap out system is a provider access system. We recorded 1 instance of 0.5-hour maintenance downtime of the tap in tap out system

CDS

A CDS (Clinical Decision Support) System helps ambulatory and inpatient care providers by providing EHR alerts, workflow pathways, up-to-date drug information and safe medication prescription. Clinical workflows allow for greater accountability and patient safety. There were 6 instances of CDS downtime in our records, all of which were scheduled for a cumulative downtime of 3.85 hours. Our CDS systems are embedded into the EHRs we use (EPIC and Cerner)

TIMELINE

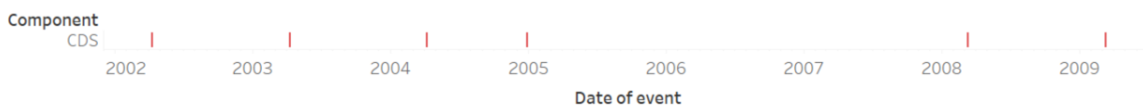

Legend : CDS – Clinical Decision Support

DISCUSSION

Event model

Infrastructure resilience is defined as the plans for the continued functioning and recovery of computing infrastructure even if some components themselves fail to work[36]. Reliability is the degree to which a computing component is free from failure. Availability is the amount of time a system is functionally usable[37]. Replicating components improves reliability but also increases the cost. A 24/7 system with 99% reliability means that it will have 87.6 hours of downtime, while a 99.9% reliability means only 8.76 hours of downtime. Considering the complex sociotechnical nature of a CCI, a failure event that causes downtime is a function of many factors.

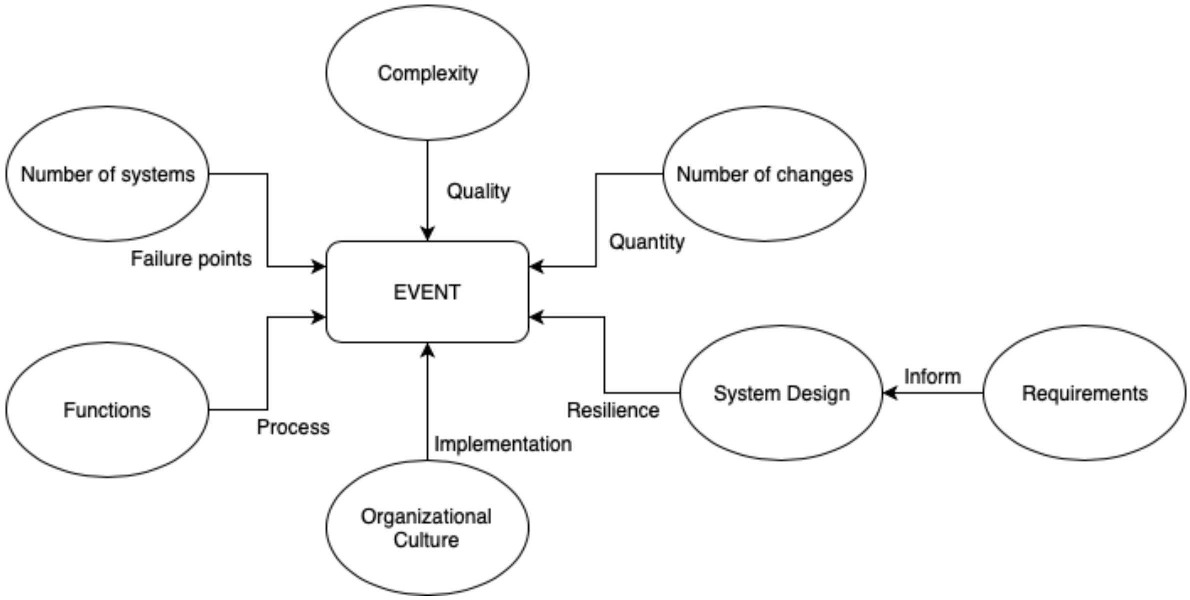

Image courtesy of Kent Beckton

Complexity of a system is determined by the specialized components that make up the system. There is a tradeoff between complexity, which rises with greater functionality. The simpler the system the lesser the chance of failure. The higher the degree of specialization, the greater the functional efficiency but also the greater is the chance of failure. The greater the functional efficiency, the greater is the user reliance on that technology. This implies that when highly specialized components fail, the greater the user impact. This has implications in resiliency planning. The greater the number of changes made to a system, the greater the number of pathways that need to be tested. This reduces system resilience. The spike in events in 2009 can be explained by the major hardware and software changes we were putting in place during this time. System changes lead to higher probabilities for downtime. The system design or architecture, which in turn is a function of requirements can have implications on the frequency of events. After the phase out of Mindscape, ORCA was implemented in Aug 2003 at our institution. With this transition came many CCI failures. Organizational culture determines the reporting and mitigation strategies institutions have in place to deal with failure events. The spike in the downtime period and the number of records in 2016 was due to an improved reporting system we put in place. The greater the number of functions and the number of components that perform these functions, the greater is their chance of failure and therefore higher the probability of a failure event. Resiliency planning means a careful consideration of the availability of 3rd party systems. As an example, our drug database is provided by a third party. Failure of this system means that our providers may not have up-to-date access of drug – drug interactions or the latest prescribing information. This could lead to sub optimal or inaccurate dosing, leading to patient harm. At UW Medicine, the goal is to push for simpler more reliable systems with redundancies planned strategically to reduce overall cost of operation.

Clinical impact of infrastructure failures

Failure of EHRs are inherently disruptive and downtime procedures (of varying efficacy and effectiveness) are usually implemented. The loss of functionality can be extreme. Clinicians may not be able to access patient records or document care delivered. Orders may not be put in or viewed. Sometimes the loss of functionality is isolated within a system, PowerChart for example, while other components may work fine.

Clinical Scenario: Citrix failure

A busy physician wants to get home to take care of errands. She wants to complete patient charts right after dinner, so her notes are in for other users. Unfortunately, she is unable to do so because of Citrix failure. She tries upgrading the Citrix client but that doesn't solve her problem. She is left with no option but to drive back to the hospital to complete her clinical notes. When she gets to the hospital, she notices that others have the same issue. In addition, all communications through CORES are dysfunctional. She is thoroughly frustrated.

Failure of LIS means that clinicians can't order or see lab reports. Failure of pathology informatics tools means that reports are incomplete with no coding information. Consults have to be obtained from pathologists, which wastes provider time and delays care delivery.

Clinical Scenario: LIS failure

It's a busy day - a building caught fire and several burn victims are rushed to the ED. This is the absolute worst time to have an LIS outage and the unthinkable happens. Physicians have no electrolyte or hematocrit information to make critical lifesaving decisions. They resort to Point of care testing until LIS systems came back online again. LIS function is restored in 2 hours, but the lab is now tied with manually back entering lab values in the downtime, creating unsafe conditions and the potential for human error.

Failure of PACS means that clinicians can't store or access clinical images, leading to delays in treatment decisions or surgical procedures. There is a potential to bypass this using the viewer built into x ray machines and angiography/fluoroscopic imaging devices. However, this approach is only reserved for scenarios that need real time images. Failure of the RIS system implies that referring physicians need to call the radiology department for imaging study interpretations. This leads to wasted time and leaves potential for medical errors[38]. Failure of Centricity web server means that clinicians cannot access patient images, causing delays in care planning and delivery.

Failure of the retail pharmacy system may mean that patients need to revisit their pharmacy for their outpatient medications. This may lead to patient drop-offs and inadequate care delivery for the patient population. Failure of the Perinatal system in the NICU setting is can have serious consequences. Delays in care in the antenatal setting can cause delays in care and increase maternal morbidity.

Failure of eSign means that clinicians can't place treatment orders electronically. They may have to resort to physically signing orders which may result in delays in carrying out orders Failure of ePrescribe means that providers will have trouble filling out online prescriptions and may need to resort to paper based methods which can be error prone. Patients may also have trouble filling out their prescriptions at their local pharmacy. Failure of AIMS may be very expensive for the hospital because surgeries may need to be rescheduled. Failure of Muse EKG or Echo means that cardiologists can't provide pre-operative evaluations or access EKG data to order or adjust medications or studies. Failure of the Pyxis system implies that nurses and other providers have to rely on manual methods to dispense medication. This could lead to errors and delays in care to get provider sign offs.

CBORD failure means that patients dietary preferences are not easy to keep track of. This may be of clinical relevance in low sodium, low potassium diets in renal failure patients. Clinicians don't have access to archived ultrasound scans when VNAs fail. Failure of electronic whiteboards implies that productivity of the clinic is disrupted often leading to unnecessary wait times, scheduling errors and billing errors adding costs to the care delivery system. Failure of the endoscopy imaging system means that imaging data was not available to plan for surgery.

Interface engine downtime can lead to insidious consequences because end users may not even know what is wrong. Unlike EHR components that are inherently disruptive, interface engine failures manifest as outdated data. All the components work well, only the users don't see the most recent information. This can lead to delays in orders and treatment plans, raising costs of care for patients, families and hospitals. There is also a possibility of patient harm if providers believe that they are seeing the most recent information and take actions based on the same. Considering the user impact of interface engine failure (under 200 users) this would be an important place to focus resiliency efforts.

> Clinical Scenario: Interface Engine failure
>
> *An attending physician needs to evaluate whether a blood transfusion is needed for a patient recovering after surgery. She reviews the test results that indicate a low hematocrit (blood count) that was obtained an hour ago and believes that it is the most recent result. Unknown to her was that the interface engine that routes lab results from the lab to EHR is having a downtime. She orders a transfusion, when it is not needed, causing the patient to receive an unnecessary blood transfusion.*

Failure of BMDI means that clinicians need to switch to manual reading and reporting of bedside monitor data, leading to errors and critical patient harm if mistakes are made. One thing to note is that due to the complex nature of the software to perform integration, most of these outages are unscheduled in this category. Clinical toolkit is often the very first page residents and other clinicians open to access other resources. Inadequate functioning may impair the clinician's ability to find the right resource. Failure of the Reg/ADT module means that patients cannot be registered, admitted, discharged or transferred. This can have multiple implications on patient care depending on the severity of the patient's condition and their living situation.

Applications would be forced to use out of date patient encounter information and location. Consider the case where a patient was moved from the ED to inpatient unit. Failure of ADT means that orders would go to ED, disrupting care there. The new patient intake has to be on paper, causing potential for error when data is back entered. Finding the patient would be difficult for providers and lead to delayed treatment if location has changed.

Failure of Autopage wastes radiologist's time and can create workflow disruptions leading to patient safety errors. In one instance Autopage script was faulty causing error messages to be sent to referring clinicians. Hospitals cannot make revenue or see delayed bill clearance when the billing system fails. Failures of MPI can have disastrous HIPAA related consequences. Patient misidentification is a serious issue[39] and it has happened once in our institution in the past. Network outages that impacted under 10 users could represent outages in a neighborhood clinic, while those under 50 could represent a department. Large outages definitely represent the entire hospital. Small network outages affecting under 10 users and under 50 users had the highest cumulative downtime. This is what we mean by insidious nature of these outages. The smaller impact (under 10 users) network outage could represent network outages at neighborhood clinics. From the timeline it might seem as if network outages are getting worse over the years (more events recently) however our expert interviews suggest that this is probably because of improved reporting and outage detection. Failure of VDI means that clinicians cannot access patient data at their workstation terminals. This may mean delays in OR time or care delivery time as providers may need to access data from elsewhere. VDI is highly sensitive to resource availability. Overburdened or failed networks quickly manifest as failures of VDI to 'paint' the screen on the workstation[40]. VDIs also depend on VPN infrastructure to allow secure access[41]. VPN failures mean that users cannot access their enterprise desktops on their devices. Failure of network has multiple clinical implications ranging from complete loss of functionality, in which clinicians cannot access medical records to limited failures representing the loss of a function such as a departmental system or an internet outage.

Clinical Scenario: Another instance of Interface Engine failure
A patient received a dose of intravenous steroids during a surgical procedure. This was documented in the anesthesia operative report. Due to human error the primary team did not sign off on the report. The Integration engine suffered an outage and the document was not visible to the primary team as a result. The patient had a fever, elevated white blood cell count and glucose (as an effect of receiving steroids) This could also be an infection, which meant that the patient was started on broad spectrum antibiotics, a treatment that was unnecessary. The blood sugar also rose to a dangerously high level requiring an ICU transfer. This could have been prevented if the team knew that the patient had received steroids intraoperatively.

Power outages and network outages have an overlap because our dataset includes network outages in UW affiliated clinics, which may not have power redundancies for routers. Failure of the firewall can manifest as data breaches or unauthorized network access. This can have significant HIPAA related consequences. The example of a firewall outage is a clear distinction where the percentage scheduled events doesn't tell us the whole story. Only by looking at the cumulative downtime we understand that unplanned firewall outages can have severe consequences (2.5 hours in under 500 users, and 3 hours in Under 1000). Failure of SAN can manifest as slow data request handling and data packet loss. This can lead to frustration among clinicians and reduce productivity by creating application slowness. Failure of the VPN system means that clinicians need to be at point of care to access patient data and document care. This can lead to delays in care delivery. Here is another example where the percentage scheduled events don't tell us the whole story. Only by looking at the cumulative downtime we understand that unplanned VPN outages can have long downtimes. Failure of the Switches, Routers or DHCP server means that devices cannot communicate, and it might manifest as network failure. Unscheduled data center downtimes are disruptive to clinical workflows since clinical communications and data transfers are impacted. Failures of the support infrastructure can cause data center issues.

Most UPS related downtimes were planned, explaining the 10 hours of user impact with zero users impacted. Failure of servers leads to a specific organizational dysfunction. Failure of computational servers leads to sluggish interfaces rendering them from frustrating to unusable. File server failure reduces the availability of critical files. Failure of SQL servers and database servers reduces data integrity and search ability. Name server failure and web server manifests as network failure. Failure of the active directory means that users cannot send emails, schedule calendar events or look up other user's calendars for availability. Pager downtime is not acceptable since mission critical communications may be lost causing patient harm and delays in patient care. More than half of our downtime events were unscheduled, which could be a cause for concern. Although we have had no pager related concerns since mid 2016, resiliency is critical for this component. Failure of the phone lines means that providers can't communicate or get a quick referral. Providers won't be able to register complaints or get teleservices. This leads to delays in care. Failure of the DAS system means that wireless services are disrupted across the enterprise. This could also manifest as a network failure. Failure of SharePoint services mean that document storage and collaboration are affected. This is particularly relevant to clinical toolkits which are hosted within our enterprise SharePoint for security reasons. Failure of PSN means that reporting safety related events is impaired.

Failure of dictation systems means that transcriptions have to be done manually and that may mean delays in notes entered into the system and delays in care because notes were not available. Failure of the intranet means that staff collaboration may be impaired. Failure of the UAA system means that providers are unable to login to systems to provide care. Lock outs cause frustrations among users. Workstation downtimes mean that clinicians have to wait to access patient information. Since most updates (maintenance related) allowed end users to restart machines at their convenience or allowed automatic updates when the end user was not around, service interruption was minimal. Failure of print servers and MFDs can result in provider frustration, particularly during rounding and billing. Failure of email is extremely frustrating to providers since it is the primary mode of communication outside CORES. Email is used to disseminate census management information at both HMC and UWMC twice daily. Inappropriate triage could result as a consequence of email failure because of the mistaken assumption that beds were available and accepting transfers. Scheduled admits may even be disorganized because of email failure. However, most email downtimes were scheduled, letting users use an alternative means of communication. Failure of tap in tap out results in delays in care because providers may be restricted in their access to care facilities. Failure of the CDS system means that clinicians may not have the latest prescribing information. Providers may override these systems but there is an audit trail of an alert and an override. Failure of CDS means that this audit is not available for root cause analysis if something unanticipated happens.

Downtime of internet explorer means that clinicians cannot access web-based clinical applications, leading to delays in care delivery.

LIMITATIONS

There are important limitations to the dataset that serves as the foundation for this report. Not all CCI failures are included in this dataset. It contains only a representative sample of clinically significant CCI failures at our institution. It doesn't contain incident reports from individual clinicians that were not serious enough to be escalated the attention of IT Services manager and leaders. Although our dataset contained a few CCI failures that did not have a clinical impact (those affecting business and research systems for example), we did not include them in our analysis, considering the clinical focus of our work. Due to the sensitive nature of the information, our data set does not include security related incidents at UW. System failures detected by backend monitoring but resolved promptly are not represented in this dataset. Our dataset is mostly human operator-reported data, filled out to the best information available and the ability of the person filling out the outage report. Some data fields are therefore missing (scheduled downtime vs unscheduled downtime). Cumulative downtime is just a proxy for user impact. A 20-minute EHR downtime during the day has more impact than a downtime in the early morning hours (2 am -6 am). Similarly, number of events is another proxy for user impact. Greater frequencies don't necessarily mean a component is failure prone. User impacts are estimates and are missing in many reports. One way to mitigate this issue is to look for machine generated system failure logs. These however do not give sufficient context to interpret the data. Usually an event contains an initial report, multiple follow up reports (particularly for longer unscheduled downtimes) and a final 'all clear' report. During the classification process, only the most recent or relevant downtime is recorded as 1 instance. Using appropriate machine learning algorithms this process can be automated, instead of having to rely on human classification, which can be tedious and error-prone. Since the emails and other reports of CCI failures were collected by Dr Thomas Payne, it is possible that not emails were received, or other forms of notification were followed to inform him. Planned downtimes can have unplanned downtimes too (longer than estimated downtime, downtime because of unexpected change) Timelines represent date of event – not date the email was sent. The log, although not complete, it could represent a reasonable estimate of the spectrum of CCI failures that can occur in an academic medical hospital. A true estimate of CCI failures is hard to estimate because:

1. Not all problems are reported. IT teams may not recognize some issues in a timely manner or have mechanisms to report them.
2. Reports may contain incomplete and inaccurate information because they don't reflect the results of a root cause analysis that may be performed later.
3. The variation in frequency of downtimes is probably because of differences in reporting. The data might falsely give the impression that the incidence of a problem is rising even if the actual incidence is declining because of increased rate of reporting. We did partly account for this in the way we filter out logs in our data processing.
4. This may not represent the experience in other institutions because the spectrum of vendor products and architecture vary widely. Many organizations remotely host their EHR (we host it locally).

Due to dependencies and the front-end nature of EHR, many downtimes may manifest as EHR failure, when in reality they are just symptomatic of the underlying issue.

Network failures are probably the most mislabeled because many component failures manifest as a network failure. One such example is when administrators shut down the network when they detect firewall failures to protect systems. So, some network failures may have manifested from firewall failures. Since Horizon Surgical Manager was integrated with Horizon Performance Manager(HPM), we have not reported data post integration considering HPM is a non-clinical performance management tool. In our dataset power related incidents were also recorded in non-data center settings as in neighborhood clinics. So, this data does not exclusively represent power issues at a data center. Since reporting thresholds vary, the data set might over represent some issues over others.

CONCLUSIONS AND FUTURE WORK

The dataset used to generate insights may be limited in its completeness and accuracy. However, there is no other data source that comes close to representing the broad range of CCI failures at a large healthcare institution. From our analysis, it seems possible that reported small-impact CCI failures (user impact under 50 users) downtime may cumulatively exceed EHR downtime. We believe that in the future, automated classification systems should be able to look at failure logs and predict failures beforehand, so system administrators can prepare for unscheduled downtimes. CCI failures at other organizations can be used to confirm and extend the findings from this study. Failure risk can be mitigated by triaging and ranking kinds of failure according to their potential for clinical impact. Redundancies can be built in for critical systems to reduce implementation costs. Ultimately, this will help advance our understanding and inform the development of safe and resilient CCI that reliably serves the needs of patients and providers.

ACKNOWLEDGEMENTS

The author would like to acknowledge the contributions of Kent Beckton for technical inputs and event models, Dr Thomas Payne for the data set and clinical inputs, Dr Mark Whipple for inputs on ideas and organization, and Dr Andrew White and Dr Fischer Lees for clinical scenarios.

BIBLOGRAPHY:

1. Hsiao, Chun-Ju, and Esther Hing. *Centers for Disease Control and Prevention*, Jan. 2014, www.cdc.gov/nchs/data/databriefs/db143.pdf. Accessed 18 May 2019.
2. Larsen, Ethan, et al. "Implications of electronic health record downtime: an analysis of patient safety event reports." *Journal of the American Medical Informatics Association*, vol. 25, no. 2, 2017, pp. 187-191.
3. Nelson, Nancy C. "Downtime procedures for a clinical information system: a critical issue." *Journal of Critical Care*, vol. 22, no. 1, 2007, pp. 45-50.
4. Jones, Spencer S., et al. "Guide to Reducing Unintended Consequences of Electronic Health Records." *ONC | Office of the National Coordinator for Health Information Technology*, Aug. 2011, www.healthit.gov/unintended-consequences/. Accessed 18 May 2019.
5. Walker, James M., et al. "Health IT Hazard Manager Beta Test." *AHRQ National Resource Center; Health Information Technology: Best Practices Transforming Quality, Safety, and Efficiency*, May 2012, healthit.ahrq.gov/sites/default/files/docs/citation/HealthITHazardManagerFinalReport.pdf. Accessed 18 May 2019.
6. McDonald, Clement J. "Computerization Can Create Safety Hazards: A Bar-Coding Near Miss." *Annals of Internal Medicine*, vol. 144, no. 7, 2006, p. 510.
7. Terry, Ken. "Cerner Hosting Outage Raises Doctor, Hospital Concerns." *InformationWeek*, 26 July 2012, www.informationweek.com/healthcare/electronic-health-records/cerner-hosting-outage-raises-doctor-hospital-concerns/d/d-id/1105525. Accessed 18 May 2019.

8. Hanuscak, Tara L., et al. "Evaluation of causes and frequency of medication errors during information technology downtime." *American Journal of Health-System Pharmacy*, vol. 66, no. 12, 2009, pp. 1119-1124.
9. Sittig, Dean F., and Hardeep Singh. "A New Socio-technical Model for Studying Health Information Technology in Complex Adaptive Healthcare Systems." *Health Informatics*, 2015, pp. 59-80.
10. National Research Council. Health IT and Patient Safety: Building Safer Systems for Better Care. Washington, DC: The National Academies Press, 2012.
11. Singh, Hardeep. "Defining Health Information Technology-Related Errors." *Electronic Health Records*, 2013, pp. 27-35.
12. "Sentinel Event Alert, Issue 42: Safely Implementing Health Information and Converging TechnologiesJoint Commission." 11 Dec. 2008, www.jointcommission.org/sentinel_event_alert_issue_42_safely_implementing_health_information_and_converging_technologies/. Accessed 18 May 2019.
13. Belmont, E. et al. "Emergency Preparedness Checklist for Information Technology Infrastructure and Software Applications." www.state.in.us/isdh/files/AHLAEmergency_Preparedness_Checklist.pdf. Accessed 18 May 2019.
14. McBiles, Mike, and Anna K. Chacko. "Coping with PACS downtime in digital radiology." *Journal of Digital Imaging*, vol. 13, no. 3, 2000, pp. 136-142.
15. Valenstein P, Treling CP, Aller RD. Laboratory computer availability: a College of American Pathologists Q-probes study of computer downtime in 422 institutions. Arch Pathol Lab Med. 1996 Jul;120(7):626-32.

16. Bittle, Mark J., et al. "Registration-Associated Patient Misidentification in an Academic Medical Center: Causes and Corrections." *The Joint Commission Journal on Quality and Patient Safety*, vol. 33, no. 1, 2007, pp. 25-33.
17. Kilbridge, Peter. "Computer Crash — Lessons from a System Failure." *New England Journal of Medicine*, vol. 348, no. 10, 2003, pp. 881-882.
18. "'Serious' Computer Crash Hits Hospital Trusts." *Mail Online*, 31 July 2006, www.dailymail.co.uk/news/article-398477/Serious-crash-hits-hospital-trusts.html. Accessed 18 May 2019.
19. Rosencrance, Linda. "Problems Abound for Kaiser E-health Records Management System." *Computerworld*, www.computerworld.com/article/2548231/problems-abound-for-kaiser-e-health-records-management-system.html. Accessed 18 May 2019.
20. Myers, R.B., et al. "Review of Reported Clinical Information System Adverse Events in US Food and Drug Administration Databases." *Applied Clinical Informatics*, vol. 02, no. 01, 2011, pp. 63-74.
21. "AHRQ 2011 Annual Conference." *AHRQ Archive--Home Page*, archive.ahrq.gov/news/events/conference/2011/index.html. Accessed 18 May 2019.
22. "SAFER Guides." *ONC | Office of the National Coordinator for Health Information Technology*, www.healthit.gov/topic/safety/safer-guides. Accessed 18 May 2019
23. Van Eaton, Erik G., et al. "Safety of Using a Computerized Rounding and Sign-Out System to Reduce Resident Duty Hours." *Academic Medicine*, vol. 85, no. 7, 2010, pp. 1189-1195.

24. Goldberg, HI, et al. "Internet access to patients' records." *The Lancet*, vol. 351, no. 9118, 1998, p. 1811.
25. Gürsel, Güney, et al. "Interoperability in Laboratory Management Information Systems." *Advances in Healthcare Information Systems and Administration*, pp. 118-135.
26. Choplin, R H, et al. "Picture archiving and communication systems: an overview." *Radiographics*, vol. 12, no. 1, 1 Jan. 1992.
27. Ehrenfeld, Jesse M., and Mohamed A. Rehman. "Anesthesia information management systems: a review of functionality and installation considerations." *Journal of Clinical Monitoring and Computing*, vol. 25, no. 1, 2010, pp. 71-79.
28. Balkman, Jason D., and Alan H. Siegel. "Autopage and the Use of Computer Scripts to Automate Microtasks." *Journal of Digital Imaging*, vol. 27, no. 4, 2014, pp. 474-478
29. Rouse, Margaret. "What is Enterprise Master Patient Index (EMPI)? - Definition from WhatIs.com." *SearchHealthIT*, searchhealthit.techtarget.com/definition/master-patient-index-MPI. Accessed 18 May 2019.
30. DeCoster, Dennis. “Project Success Is the Sum of All Parts: in a Data Center, Uptime Is More than Just the Sum of Its Pieces.(2 More Minutes).” *Consulting Specifying Engineer*, vol. 49, no. 8, 2012, p. 64.
31. Bacik, Sandy. "Tape Backup Considerations." *Information Security Management Handbook, Sixth Edition, Volume 5*, 2011, pp. 423-428.

32. Benson, Theophilus, et al. "Understanding data center traffic characteristics." *ACM SIGCOMM Computer Communication Review*, vol. 40, no. 1, 2010, p. 92.
33. Course, Microsoft O. *98-365: Windows Server Administration Fundamentals, B&N Renting e-Bk*. John Wiley and Sons, 2011.
34. Mehrzad, Raman, and Michael Barza. "Are physician pagers an outmoded technology?" *Technology and Health Care*, vol. 23, no. 3, 2015, pp. 233-241.
35. Saleh, A.A.M., et al. "Distributed Antennas for Indoor Radio Communications." IEEE Transactions on Communications, vol. 35, no. 12, 1987, pp. 1245-1251.
36. Jackson, Scott. "The Principles of Infrastructure Resilience." *Domestic Preparedness*, www.domesticpreparedness.com/resilience/the-principles-of-infrastructure-resilience/. Accessed 18 May 2019.
37. Payne, Thomas. *Practical Guide to Clinical Computing Systems: Design, Operations, and Infrastructure*. Academic P, 2014.
38. Haux, Reinhold, et al. *Strategic Information Management in Hospitals: An Introduction to Hospital Information Systems*. Springer Science & Business Media, 2004.
39. Just, Beth Haenke et al. “Why Patient Matching Is a Challenge: Research on Master Patient Index (MPI) Data Discrepancies in Key Identifying Fields.” *Perspectives in health information management* vol. 13,Spring 1e. 1 Apr. 2016
40. Shields, Greg. "Five Ways Your VDI Implementation Can Fail." *SearchITChannel*, searchitchannel.techtarget.com/tip/Five-ways-your-VDI-implementation-can-fail. Accessed 18 May 2019.

41. Vaughan-Nichols, Steven. "The Trouble with VDI." *ZDNet*, 1 Aug. 2013, www.zdnet.com/article/the-trouble-with-vdi/. Accessed 18 May 2019.