CHAPTER 1

Healthcare IT and the Growing Need for AI Operations

Shall we play a game?
—**Joshua (from *WarGames*)**

COPYRIGHTED MATERIAL

In today's ever-changing business model of do it faster, do it better, and do it without flaws, there needs to be a balance between those who create the technology and technology having a mind of its own. As in the famous quote "Shall we play a game?" healthcare operations is anything but a game and lives hang in the balance. In today's organizations, that balance is being established as technologies such as artificial intelligence (AI) are being implemented. There needs to be a way to do more work efficiently and with greater intelligence while still ensuring that the work is performed correctly. With the boom of healthcare advancements and the need to keep up with technological change, those who rely on all of the newest technology for clinical operations need an enterprise system that ensures that the technology continues to work for us and not against us. That combination of technology and clinical advancement comes in the form of a successful merging of intelligence and strategy, using the correct tools for the job, and planning and designing a platform that works for you, not against you. This is AI operations (AIOps) in healthcare.

This chapter explores the healthcare market and how technology continually changes it, specifically within the realm of AIOps. In these pages I will discuss the growing need for technology in this space, how healthcare has been fundamentally (and forever) changed by the digital

landscape, and all of the specifics revolving around AIOps. This includes how AIOps is being used to create efficiency, reduce downtime, increase time to respond to issues, improve the ability to automate efforts to reduce waste and time spent doing computation work, and ultimately create better customer experiences for all patients, clinicians, and everyone involved in the healthcare space.

In the first portion of this chapter, I will cover the basic history of artificial intelligence (AI) and machine learning (ML). Although some could say we have always been in a perpetual state of "machine learning" for as long as we have had machines and in a constant state of computational (or artificial) intelligence as long as we could compute things, there are some significant milestones in the ML and AI timeline. For one, as long as we have been playing games, there has always been a study of game theory and outcomes through games. Many military and war strategists believed in game theory, and this became even more apparent when IBM began testing ML theory with gaming to produce the first machine learning game in the 1950s when someone played checkers and the program was able to learn from the outcomes of the game, the players' choices, and so on. I think this real story was likely the predecessor to the movie *WarGames* decades later. Checkers, chess, backgammon, and other games were all tested to see how a machine could learn.

As more and more technology (machines) was created and advanced, the same questions and theories were applied to it. When cars were made, how could we get them to learn? What about if we made a robot? Could it learn? The same theories from a long time ago all waited until technology caught up and provided for computers, robotics, and other major technological advancements that could be fused with machines to allow them to learn. Once computers were added to cars, then cars could start to learn. Now we drive in cars that can predict a possible crash and take action. This development went way beyond the abilities of game theory, but it should be noted that the mathematical equations, usage, and logic behind it still remained the same. It was only advancing as quickly as the technology did and was expanded on.

Another major installment of ML and AI development came with the World Wide Web (WWW), the Internet, and the Internet of Things (IoT), where the interconnected nature and development of all of technology was able to fuse and share data as well as save it. The saving of large quantities of data (or big data theory) allowed for more math to be applied for machine learning capabilities. Also, the growth of large-scale search engines (like Google) continued to allow for even more ML and AI abilities due to the analytics that could be applied to "customize" an

experience for every user. Augmented reality (AR), wearable technology, DNA collection, mobile technologies, social media, and so many other advancements bring us to a stage in life where ML and AI are able to be used not only in any one of these advancement areas but also across them as they too interconnect.

This is where we begin our journey into the development and fundamental layout of AI, ML, and AIOps in the healthcare world. Healthcare is the largest user of all the technologies I just mentioned and the biggest connector of intelligence usage to increase the use of treatment, medicine, and patient satisfaction. We now need to know how to keep all of these systems running and available, continue to perform their computations, and allow for the continued interconnection and learning to provide the best care possible now and in the future.

A Brief History of AI and Healthcare

No industry is bigger, more important, more dynamic, and exploding with change than healthcare. Some may argue that it is similar to the industrial revolution in its transformative scope. The electronic medical record (EMR) and other technological advancements are changing the way we see, deliver, and expect to receive our healthcare. One of the most interesting things about the healthcare industry is we are all our own clients, customers, and patients, which makes our industry unique in that it is something we hope we never have to use but absolutely must have in the form of benefits to cover our families and ourselves. Working in this field can be the most rewarding experience one can have, and seeing its growth and being a part of its transformation can be a once-in-a-lifetime experience.

As the healthcare field grows in every aspect, we must consider the technology used to bolster this revolutionary expansion. We call this technological field *healthcare information technology* (also known as healthcare IT or HIT). In this chapter, I will explain healthcare IT, some of its history, and why technology has expanded it exponentially. I will also start to talk about the need for another popular and growing technological advancement called *artificial intelligence* (AI). Beyond that, we will bridge the two technologies—healthcare IT and AI—and explore a third topic called *healthcare operations* so we can create a fusion between them all, which is known today as *artificial intelligence operations* (AIOps). Let's begin our journey in this chapter and this book by starting from the beginning, which is the expansion of healthcare as we know it today.

THE CORONAVIRUS AND COVID-19

When the coronavirus (COVID-19) pandemic spread around the world in 2020, it affected the world of medicine in a dramatic way. As I will cover in Chapter 8, "The Future of Healthcare AI," there have been radical changes to the use, delivery, and expectations of healthcare service.

With guidance by the World Health Organization (WHO), Centers for Disease Control and Prevention (CDC), and local, state, and government leaders, the world population had to make radical changes to slow the spread of COVID-19. For one, all healthcare systems (hospitals, practices, etc.) needed to find ways to reconfigure to handle the surge into the emergency department (ED) and intensive care unit (ICU) areas of the system. Temporary facilities were set up, and many new ways of practicing emergency medicine were considered. New medications, new treatments, and a world of new research into finding a cure were put in place. The world actively worked to slow the spread and focus on a cure. New technologies emerged, the importance of keeping systems up and running found a renewed priority level, and the use of older technologies saw a resurgence, like telemedicine.

Healthcare IT Expansion and Growth

The radical expansion and growth of IT and healthcare IT was just the beginning. It provided the needed building blocks to get to where we are today so that we can collect data in large amounts (*big data*), analyze it, and make predictions and assessments to create better outcomes. Big data analysis helps our knowledge and handling of population health, the need to reduce hospital stays, what can be done inside an acute versus nonacute facility, and how to make predictions on outcomes in a geographic area. An example of a prediction in an area would be how seasonal flu may impact certain areas and why that may be. This only scratches the surface of what we can leverage big data for.

To get to this point, we needed to get all data into the computer systems. By creating the electronic medical record (EMR) and having clinical staff adding this data to these systems, the building block was in place to start expanding this practice across all health systems. This provided many benefits, one of which was leveraging the resources of the many instead of the few. The increasing cost of healthcare put a lot of stress on smaller hospitals that could no longer afford to continue to build technical systems while still expanding their operations outside of information technology. Other important factors include (but are not limited to) security, risk, and compliance where privacy became paramount by law.

Compliance, meaningful use, regulatory bodies, inspections, laws and even the dominance of the Health Insurance Portability and

Accountability Act (HIPAA) passed by Congress in 1996 continued to drive more and more healthcare providers to join forces with others to share resources so they could stay afloat. One of the many benefits was the ability to leverage the administrative functions (like technology) between providers, clinicians, healthcare facilities, practices, hospitals, and even insurance. Another advantage was the chance to scale up on all of these shared resources, which allowed hospitals and practices to share operational data so that they could model best practices and standards to keep all of these technology systems operational, resilient, and well-positioned for future innovation.

BIG DATA AND ITS IMPORTANCE

Data is the fundamental building block to everything that we do in technology. Think of a simple, traditional network. A network is useless unless you have something to share. Would you spend all of this money to set up countless connections simply for the sake of having connections? Of course not. You need to share data from printers, faxes, email, text messages, and files.

Big data is the same building block for AI, ML, and AIOps. Without data, what are you collecting, and why would you need it in the first place? The key to doing mathematical and computational analysis is to find ways to solve equations, and when you apply large volumes of data and a plan for finding trends, you wind up with predictive analytics. Analytics is what you do to mine the data for whatever specific information you need. So, if you are trying to see where there are systematic flu outbreaks in an area and why the flu outbreaks are occurring there, you need to collect data, analyze it, and create a mathematical prediction based on the trends. Big data is the large-scale collection of data that would need to be mined to identify trends when you conduct your analysis.

Data Overload

The many sources of data and the need for it to be shared started to become overwhelming. IT environments continued to grow exponentially, outpacing the underlying resources needed to run them. For example, extremely large enterprise storage system platforms with terabytes of space and high-speed delivery became essential to keep up with the needs for storing and leveraging data, particularly for the data scientists who looked at, analyzed, and computed this data (big data) for many use cases.

Data overload is a potential issue. You can collect a ton of data and not know what to do with it because you lack the strategy, tools, understanding, applications, services, or staffing needed to unlock what is within it. You may also run into issues where you cannot collect, store, mine, or access in a reliable way the large data sets you collect. You may also run into issues with sharing these large data sets. For example, I have seen in the past where research teams conducting important studies could not share files with other research teams remotely because of their large size. What people may not know is that (as an example) a magnetic resonance imaging (MRI) file with large slicing capability could be terabytes in size depending on the resolution of the file in use. When sharing hundreds of these over a typical 10 megabit connection with virtual private network (VPN) encryption, you could expect issues!

Unorganized data is another potential issue. When you work with data, you need to tag it correctly so that it can be accessed via field searches in databases. If you do not tag it correctly, when you run scans to find keywords or other criteria, you may not find what you are looking for, or worse, you may miss large portions of data, which can skew or create kurtosis of your results.

Unusable data is a related concern, meaning the data is missing entirely (orphaned), misunderstood (outliers), or corrupted and therefore unusable.

In any of these cases, you can have a data grooming exercise conducted to make sure that a sample set of the data shows that these problems may or may not exist and be fixed if possible. This is generally part of an organization's master data management (MDM) program. Regardless of what state your data is in, you need to recognize that it could be problematic and cause issues in your deployments, use of AI and ML, and other future endeavors of innovation. One of the most important things to consider when conducting any mathematical or computational exercise is that your data is usable and trustworthy.

CLEANING UP YOUR DATA WITH MDM

Data is the fundamental building block to everything that we do in technology. The best way to handle your data is to start fresh and make sure that your organization utilizes a master data management (MDM) program. An MDM is more of a governance committee made up of both technology and business leaders that puts rules into effect on how all company data will be tagged, listed, used, controlled, stored, and named (conventions). This

includes identifying who has access, who owns what data, who maintains the data, and who ensures the lifelong accuracy and accountability of all data shared within the company. Without a program of this kind, you will run into issues later when your data grows so huge that it will take large project-like efforts and create massive disruption to address these and other issues. If you plan on working with data (which is the building block of AI and ML and thus AIOps), make sure that you know what the state of your data is before you begin innovative projects that require the use of big data.

Digital Transformation of Healthcare

The digital transformation of healthcare takes two paths: the clinical and the operational. Although I will cover them both in this chapter, the focus of the book really only goes deeply into one of these two paths. Although they are intertwined, they are in fact different when it comes to their stated purpose and scope. Make no mistake, they are both connected, and one may argue that you can't have one without the other. Case in point, if you do not have a stable environment in which to work (AIOps), how can you have a platform to conduct informatics to prevent illness or disease? On the contrary, why would you have a large AIOps platform if you weren't doing this type of the work in the first place? You can also argue that without the development and growth of AI and ML over the years, the fusion of this technology into newer tools like event and fault management platforms, service management tools, and help desk systems would be pointless when it comes to building on AI and its benefits. As you can see, they are both interdependent.

To further explain these two paths, path 1 is the clinical informatics space where AI and ML are used to do informatics work. Big data and its dissection are the future of clinical healthcare. As we collect more and more data, successfully analyzing it allows us to predict behaviors, solve clinical problems, create cures, or create new treatment plans. Looking at this data can tell you historically how many times the flu came through an area and at what impact rate so you can make sure that you market the flu shot effectively or logistically ensure that you have sufficient vaccines in that area. The use of this data is important to innovation and research for moving healthcare forward into the future.

NOTE **For the clinical side of AI, ML, informatics, and using data to create positive outcomes, there is a natural fear that the storage and usage of this data would somehow create a privacy concern. However, all healthcare**

systems, providers, and users of this data are mandated to only use copies of the data that have been de-identified for use. This means all identifying data that can link any medical data to a patient by name or any other form of ID is removed completely before any data is used to study. This allows for privacy to remain intact.

Path 2 is the AIOps path where AI and ML are the focus of using enterprise platform systems that keep all of the data up and running alongside the EMR and every other clinical system, application, program, and system in use in the health provider environment.

For example, consider a patient needing to go in for an annual checkup, complete with routine blood work and urine sample. The doctor reviews the results of blood and urine tests to give the patient an assessment of the current state of their health. The doctor may also look at the patient's chart to see what they are predisposed to, their family history, their age, and many other factors before giving advice, a clinical path to follow, medications, or a referral to a specialist for further work. The patient's chart is reviewed to see what their last checkup or blood labs showed, and trends and patterns are examined so that nothing is missed. Various systems support the bloodwork lab, the urinalysis lab, the pharmacies, the specialists, the patient's historic files, the patient's current general condition, the ability to recognize patterns and predispositions, and the ability to project potential areas of concern, and AI and ML play a role in synchronizing and coordinating all of this data.

One of the reasons why understanding this is so important is that it sets the basis for keeping all of these operations up, running, stable, and operational. All of this technology must be highly available, recoverable from outage or disaster, and manageable. This can be done through traditional methods, but it can also be done with artificial intelligence. There are pros and cons to both, and we will explore these considerations throughout the book. Remember, the reason we want to design and set up an enterprise system like AIOps is so you, your family, your community, and your world can get the quality healthcare you expect and, in some cases, demand. If any of these systems are down, not operational, or corrupted, you can get delayed care, no care, or, in some rare cases, the wrong care, which can bring about life-altering experiences.

WARNING **When dealing with a pandemic like COVID-19, the stability of infrastructure, the integrity of systems, and the reliability of key applications are beyond mission critical. They are always a priority, but with the need to**

keep healthcare workers focused on providing clinical care during such a dire crisis, the use of AIOps can be instrumental if deployed correctly. You need to strategize, configure, plan, deploy, and use AIOps to achieve that goal.

The Science of Healthcare Innovation

So, with all of this technology, what are we really striving to do? Much like Maslow's hierarchy of needs, we want to reach enlightenment and integrity. Once we have learned how to survive and take care of the fundamentals such as stability and integrity, we ultimately want to innovate. Innovation is the hallmark of any civilization that has moved to the highest stages of self-enlightenment.

With healthcare, the goal will be to have healthcare IT systems that run 24/7/365 with a five nines (99.999%) uptime and immediate resilience to any disaster whether technological, weather-based, military, or biological. We also want to reduce cost, waste, our environmental footprint, and our need to react to system issues. Lastly, we want to create efficiency, innovation, and the ability to rely on the systems we deploy and use. That's a lot to ask for, right? Well, this is the underlying goal that has spurred the entire market called AIOps. AIOps provides the ability to do everything I just said and more. With AIOps, your key systems will remain stable, and intelligent decisions will be made through machine learning to reduce impact related to downtime. Downtime impacts your ability to innovate.

To have innovation, we need our experts focused on their jobs. They should not be waiting for systems to recover from outages. They should not be working on how to connect users to systems as they recover. Our clinical experts shouldn't be using pen and paper during an outage and then afterward painstakingly entering that information to the EMR without a single data error. All of this is a waste of time and energy that creates missed opportunities for greatness. It also increases the potential for error.

Can we remove all of this waste and create innovation that we can all benefit from? The emergence of a technology platform that delivers AIOps promises to fulfill this role, among other things. Does it come with its own issues? Yes, and I will break them down and how to overcome them throughout this book. Regardless, the biggest goal of AIOps today is to create the ability to not have to manage failing systems and give that time back to those who would use it to create magic like innovative new healthcare solutions.

NOTE **There are many roles an IT expert can play in healthcare IT such as analyst, data scientist, technical professional, systems engineer, risk manager, and so on. Everyone in IT should be enjoined in AIOps, and all parties should be stakeholders to any project where AIOps is deployed and used in an enterprise. Others who are to be involved would be any clinicians, providers, or other staff who will use the systems. Including their voices helps ensure that systems will be developed and implemented reliably in your enterprise.**

Artificial Intelligence in Healthcare

Before we can really delve into the platforms that deliver AIOps, other supporting platforms, offerings, designs, strategies, and use cases, you should have a firm understanding of the fundamental concepts of artificial intelligence, machine learning, operations research (OR), and other technologies that I will refer to throughout this book.

First, artificial intelligence is exactly how it sounds. You have a computer system (or form of technology) that is able to learn, grow smarter, or be self-aware. It has "intelligence" between its programming language and expected behavior. This is a far cry from the Cyberdyne Systems' Skynet in the *Terminator* movies. Will we one day have the type of adaptable, self-aware, artificial intelligence technology system like we see in science fiction? We don't know. However, for here and now, the answer is no, we have nothing like that. Today's AI is not Skynet. AI today is still in its infancy, and while it is mature in the sense that we have adapted it into many of our technology systems, the term *artificial intelligence* is more of a play on words than a truly conscious technology system.

In simple terms, current AI can be thought of as computer technology such as programming languages, applications, and systems that emphasize fabricated, simulated intelligence through patterns and trends. It is not conscious, but it is programmed to be intelligent. I think this is an important distinction to make right up front to make sure that you are working from the beginning with the correct definition and expectations.

As we develop AI systems, we can think of them as behaving as if they have consciousness. In doing so, it can help to further divide AI into hard AI and soft AI (also known as strong AI and weak AI). Soft AI is simple. It is usually purpose-built and performs a single or simple function. This can be thought of as a program that you operate such as a game that tracks and learns your skills, patterns, and game-playing behaviors. Hard AI is something more complex with deeper

programming, abilities, functions, and architecture. A great example of hard AI would be Google Brain. Google Brain is a hard-AI, deep-learning, AI-based research program formed in the early 2010s by Google that incorporates the concepts of machine learning with the large-scale computing resources of an enterprise as large as Google. Hard AI such as this can be more successful at simulating intelligence because of the complexity and power scaled behind it.

Either way, whether hard or soft, AI is still something that needs to be developed, worked on, built, deployed, and managed, and once everything starts to really work well, it needs to be managed some more. This is an important point to keep in mind as you embark on your journey into the world of AI in healthcare. AI requires real time and effort. So, now that we have covered AI, how does machine learning fit in?

Machine learning is the science behind the AI machine. ML can be thought of as the mathematical algorithms, statistics, underlying models, and data that powers AI. ML can be seen as a subset or component of AI where AI is what is used to make decisions and ML is what gives AI the ability to learn new things to make decisions about. ML has many components and relates to many topics such as data mining, analytics, informatics, and AI. An in-depth discussion of these topics is outside the scope of this book, and each of them can be adapted into a book of its own. I will, however, point out how healthcare AIOps and ML cross-connect and that is that there is a need for ML to be the underlying technology in AIOps for the technology to work correctly.

NOTE **There is another concept to note, which is operations research (OR). OR is a "decision-tree" concept of AI where ML is the underlying science to create data for decision making, AI is the system that allows the decisions to be made in an automated way, and OR is considered the overlay to contrast with AI in that it is the decisions that are made with or without the automation. The "intelligence" part comes from the automation of actions, which we will cover in depth when we talk about the enterprise systems that allow AIOps to become a reality.**

Now that you have a good understanding of AI, let's fuse this concept to healthcare. Healthcare is an ever-changing, evolving, dynamic field with many components. I like to divide healthcare into two separate realms when discussing AI: clinical and operational. These concepts were introduced earlier in the "Digital Transformation of Healthcare" section, and we will explore them further here.

Clinical AI is the use of AI, ML, and other intelligence components for clinical outcomes. As an example, if I want to trend what areas in a state are highest for the flu each year so I know where to stockpile medicine, I can data mine that information and make a clinical decision based on the results. This is the part of AI where ML allows us to use data mathematically to make decisions. Within the clinical realm, this use of AI and ML is still in its infancy, but major efforts are under way to develop innovative new technologies, tools, programs, systems, and workflows to more efficiently and effectively use clinical AI. Since this book focuses on the "operational" artificial intelligence realm, I will not delve too deeply into the clinical side except to make comparisons or show how clinical and operational AI at times will cross paths. However, it is important to know that AI in healthcare is very much a budding field where the positive clinical outcomes we seek and the innovation that is required to realize them is only just emerging.

AI in healthcare in the operational realm, or AIOps, is where most of our discussion in this book will focus. Operations can be defined as functions that keep the day-to-day business moving forward and ongoing. You can think of installing a server as a project where the management, maintenance, and patching of that server is the operational component needed to ensure the integrity of service. Not just healthcare, but all verticals, channels, fields, and roles require an operational component. Figure 1.1 shows the build-out of design, build, and run, which is a model that is commonly used to describe how we normally run an operation.

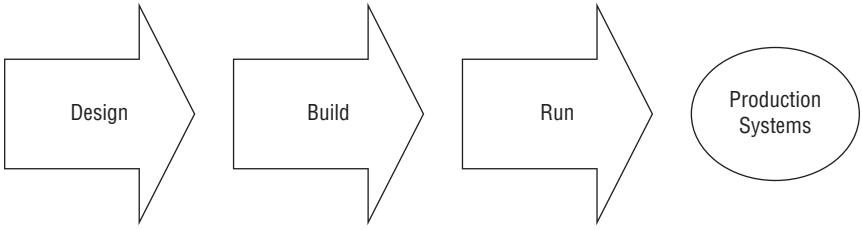

Figure 1.1: Design, build, and run model

The design, build, and run model is used when we want to do something like deploy a service (for example, a medical service like a new EMR system). First we design the service, then we build it, and then we run it.

The design phase is where we put together a strategy, create a project plan, procure funds, put a team together, and so on. The build phase is considered to be an engineering component to the workflow. We make something that needs to be managed or a service that needs to be maintained. At the end of the build process, the service is tested and, assuming it functions properly, is then signed off as production ready. Next, we run it. This is the operational component to the workflow. Now, it enters into monitoring systems so that it can be safeguarded from failure.

This is normal operational behavior for designing, building, and running anything, and in the healthcare space it is no different. However, there is one major difference to note that is a priority. The biggest difference that makes healthcare IT, healthcare operations, and other functions of healthcare service delivery more important than most other fields is that human lives are on the line. It is critical to keep in mind that there is a real "life or death" aspect to the things we do in healthcare. Therefore, when we monitor systems for use in a healthcare setting, having them operational and ready for use is an absolute top priority. This is where AIOps meets healthcare operations and helps to not only bolster that priority but act on it as designed.

USING ITIL IN HEALTHCARE OPERATIONS AND AIOps

Although this book is specifically about healthcare AIOps, being aware of how operations are generally built for the AIOps platform will help when you design and deploy your operations. One of the most widely used models that establishes IT best practices and guidance that you can use to deploy tools correctly is the information technology infrastructure library (ITIL). ITIL is a model used to help create a sound and reliable IT organization and is an important framework for modeling how to deploy reliable IT services. For example, there are five major sections to this model, which include service strategy, service design, service transition, service operations, and continuous process improvement. Each component of ITIL digs deeply into how a service is created all the way from inception through transmission and into production. It also creates an iterative function with continuous service review so that things can be improved especially as organizations and their services change. You can see an example of this in Figure 1.2.

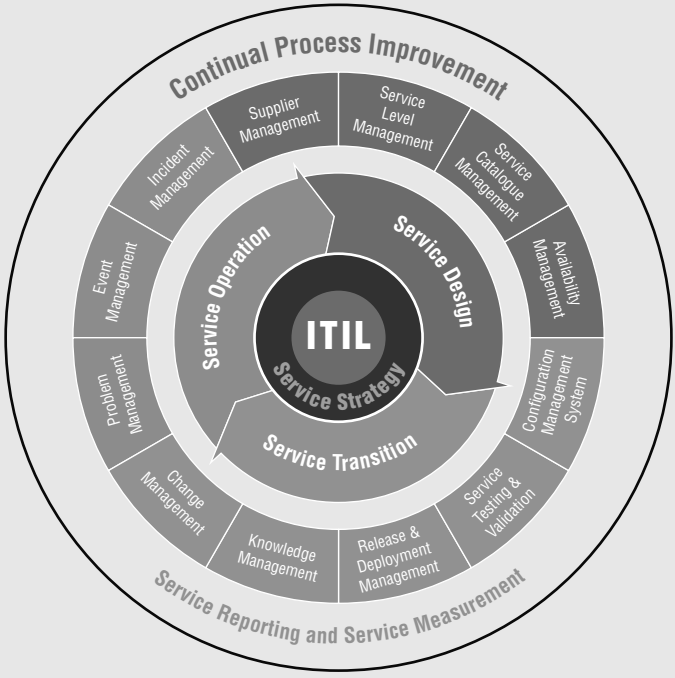

Figure 1.2: The modules of ITIL

The important takeaway here is that, as I noted with the design, build, and run model, services and systems need to be created well and managed in an ongoing and proper way so that they provide the required services. In healthcare operations, this becomes the fundamental factor of importance. When deploying AIOps, frameworks such as ITIL can help to ensure that once you get your services designed and deployed correctly, AIOps can deliver real value as an overlay and not just add confusion to poorly managed services that need some attention and focus.

Keep in mind that you can't build a good house on a weak foundation. Make sure that you consider the design suggestions listed in this section prior to your AIOps deployment to ensure that the house you build will be stronger and last longer.

Healthcare IT Operations

So, how does all of this fuse together to connect technology, healthcare, operations, and artificial intelligence? When you have a platform that can predict and respond, you have reached enlightenment. Today's operational systems usually lack interconnectivity and the ability to

leverage the benefits of convergence. If we want to successfully navigate the roads of healthcare IT, we must build and pave them correctly. Only through proper planning and implementation will we deliver AI into healthcare IT operations.

The use of IT in healthcare is already deeply rooted and will continue to grow. Working in healthcare with paper files has been replaced with the EMR system and its interconnected world of computer technology that provides a clinical facility or doctor's practice more insight into how to provide care. Keeping that functionality up and running at all times can be complex and present an array of challenges. For example, if interconnectivity between systems is not operable, you may be missing components of the healthcare service you are trying to deliver. Let's look at some real-world scenarios where this might happen. If you want to get the results of bloodwork, you need the connection to the lab that processes that information, possibly with the proper credentials for security purposes. In another instance, if you want to prescribe medications, you need connections to the patient's preferred pharmacy (or the proper department if you're providing them within a hospital or other medical facility). In prescribing these medications, if your system has artificial intelligence that can check the patient's current medications and flag potential contraindications between them and your new prescription, that would be very helpful. To make such clinical decisions successfully and to the greatest benefit possible, you need the full, accurate data and functionality of all parts of the system.

Healthcare IT is the fundamental bedrock of allowing clinical services to operate with the technology that it deeply relies on in this day and age. HIT also contains elements of risk and the need for security, which is a critical part of HIT. That acknowledged, now let's consider how AI relates to HIT. AIOps is simply the fusion of HIT with AI that enables an enterprise system or service to take the data (such as that provided from ML and data mining) from production systems, learn patterns about ongoing production and day-to-day operations, and make predictions and decisions based on that data. For AIOps to function optimally in doing this, it needs to be mature, set up correctly, and given enough time to perform its tasks.

Applying AI to IT operations is one of the biggest benefits to healthcare operations as a whole. If we can deploy a system that can find our weak spots and either self-heal them or at minimum alert us to failure and make educated suggestions (or decisions) based on reliable data, we have evolved into the realm of enlightenment. This is one of the primary goals of all AIOps platforms. There is much that needs to be considered

before we can realize that goal, and we will explore these considerations throughout this book. For the moment, I want to address fear.

There are many fears about this type of technology, and I will work to dispel any concerns with real-world experience on the subject whenever possible or share what other experts have uncovered in their strategizing and deployment of AIOps platforms in their environments. Most healthcare professionals' biggest fear would probably be, "If I deploy an AI solution that makes decisions without my approval, how would I know that those decisions are the right ones? What if a mistake is made?" Well, that is why this book will carefully examine all deployments of AIOps and how you can ensure that you are working with the best possible outcomes prior to and after deployment of these systems.

AI AND AIOps IN USE TODAY

Although this book will focus on enterprise platforms that deliver AI to manage large-scale operations, you should be aware of how hard and soft AI is used today all around you in healthcare, in IT, and in general. Today, everybody is delivering AI or AIOps. Every enterprise tool, platform, system, or application has the buzzword "AI" attached to it. The reason, quite simply, is that it sounds cool. More importantly, if we can achieve true AI, the benefits could be beneficial to us beyond our wildest dreams. But even the current state of AI is impressive. It is also pervasive. There are AI platforms in use today everywhere you look. Your mobile device, TV, and smart home device can connect to each other to use AI even more powerfully. Applications you use every day are AI functional. Obviously, the biggest ones are intelligent and can use functions like speech recognition such as Apple's SIRI, Google's Now/Home/Brain, Amazon's Alexa, Microsoft's Cortana, and others. Problem solving through decision trees can also be considered AI. Other platforms such as security systems, camera systems, and surveillance and monitoring systems can be automated to trigger alarms based on facial or other biometric information. There are so many others, too. Be aware that all of these functions have found their way into clinical operations and innovations as well as the HIT arena.

Healthcare AIOps functions can include technology we have already seen for years and is evolving into more AI-enabled platforms. These include IT technology such as antivirus (AV) platforms that use heuristics to make decisions on malicious software (malware) that can infiltrate your network and systems. For example, your AV platform can be loaded with a definition file so that it can analyze data and look for anomalies. This is one of the oldest forms of information technology and

AI from decades ago where you believed that your computer was smart because it could identify an attack pattern from a virus and quarantine it for you. The truth is, you had to keep those systems updated every single day (and you still do) with definitions that point the engines of the application to what those anomalies should look like. The heuristics function is the examination of that engine and the data that passes through it to make a "decision" or conduct an action based on that data. In mathematical terms, the heuristic function is making an intelligent decision based on information it knows to be right or wrong based on a threshold. As an example, an AV system might know that if it sees a high enough amount of a particular pattern on the system, that may be the sign of an attack, so the AV might send an alert to those who manage the system, flag the file as potential malware, or quarantine the file. This form of AI is the precursor for today's automated AI functions in a much larger world.

Another example of this type of technology is email spam filtering (based on rules), where if you flag something as spam, your email system will look at your email for patterns using data such as keywords, sender, Internet Protocol (IP) address, hostname, Domain Name System (DNS) server, and other technical qualifiers and automatically isolate those types of emails so you do not open them and spread a virus.

The same technology that local systems use for AV and spam filters can be seen in enterprise tool systems such as network-based intrusion detection and prevention systems (IDPSs) where the system captures data that traverses a network, recognizes potentially malicious anomalies, and blocks connections or quarantines data. The quarantining and blocking actions are considered the "intelligence" portion of the equation. However, you must remember that this information is acquired, learned, and ultimately refined at some point to identify outliers that do not match the proper patterns and might cause issues to the system if undetected. For example, you want to know if an email was erroneously flagged as spam and then quarantined when it actually should have been delivered. This requires tuning, refinement, and relearning. As we can see from systems of old, mistakes can be made just like in any other "intelligent" creature we know.

Figure 1.3 demonstrates this technology that AV systems, email spam filters, and IDPSs share. Basically, a heuristics engine keeps and updates a list of definition files that it can use to flag potential malware, scans data input for potential malware, and places malware in quarantine. This is one of the most common uses of AI in technology and the beginning of AIOps with security systems.

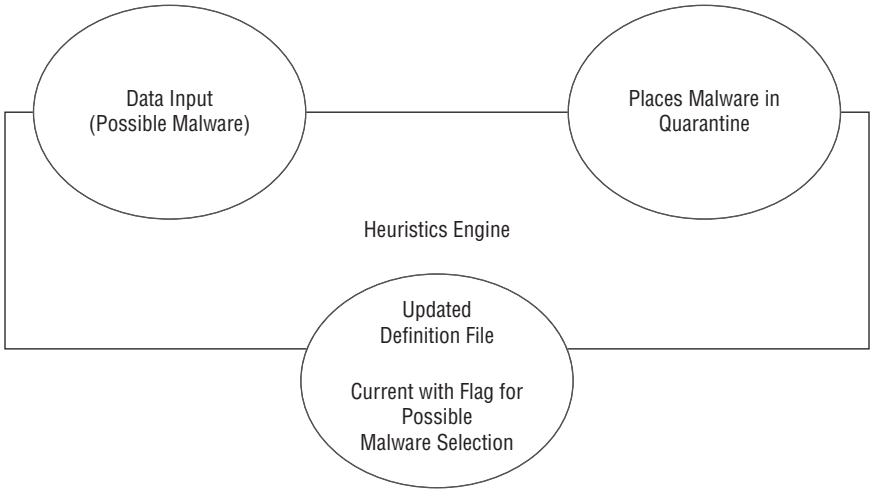

Figure 1.3: A heuristics engine using AI

Anything automated can be elevated to use AI. What truly makes it smart is that it has been configured to act that way. What makes it autonomous is its ability to do this without being told. You want to help your system become so smart that it borders on being autonomous so you can unlock its full power. Throughout this book, you'll see how to make smart, aware, and strategic decisions that allow you to benefit from this technology in the real world. When it comes to healthcare delivery, there can be no mistakes. Next we move into today's enterprise platforms that supply the AIOps experience to widescale systems and in a healthcare environment.

AIOps Platform Strategy

The impact of AI in health systems, EMRs, and other toolsets cannot be emphasized enough. Many AIOps platforms exist, but we will focus on the ones directly related to healthcare IT and how they can be leveraged to provide benefits to your enterprise architectures. Although we cannot cover them all, we will cover the most popular offerings. Before deploying an AIOps platform, your strategy has to be well thought out and executed with purpose. Make no mistake, purchasing this technology is expensive and will require approvals for budget dollars that may be questioned, so you need to select a great tool from a company that can back it with support, guidance, and documentation. You may require consulting fees

and education credits for ongoing learning about the tool. You may need to educate a portion of your staff to operate and maintain the tools you purchase. Once you have selected a platform type, you will then need to know the players, options, services, functions, and offerings within that platform, and knowing the differences can make all the difference in your selection. Let's take a look at what those types are and how you can make a more informed decision on what you need for your strategy.

Platform Types

When considering AIOps platforms, it's important to know that just because something is called AIOps doesn't mean it's a plug-and-play tool that automatically learns everything on your network and makes great decisions. Quite the contrary. This is a complex environment that requires an understanding of the inner workings of service delivery, infrastructure, your core business, and all of its meticulous underpinnings. That said, let's break down some of these offerings so you are aware of how they differ and how they interconnect. In particular, we'll examine enterprise monitoring, information technology operations management (ITOM), and application performance monitoring (APM).

The first platform we'll examine is the enterprise monitoring and AIOps space. Tools are deployed enterprise-wide to encompass all architecture so that operational data can be ingested and learned and automation functions can be selected and applied based on patterns and trends that need to be considered. For example, Splunk (`splunk.com`) is an industry leader for AIOps and has been in business developing mature solutions for more than a decade. I will discuss Splunk further throughout this book. This type of platform can be on-site or cloud-based. The Splunk service can ingest data from connected systems and allow for high-level automated actions to take place. You can set up your network, servers, services, applications, systems, and other infrastructure on Splunk, and it will start to cross-examine the data to understand patterns between them to help you make educated decisions. When fully matured, Splunk can even be automated to make those decisions for you. An example may be that all of your network switches are performing poorly at and around 9 a.m. every workday for about 20 minutes. This trend is found to be caused by an excessive amount of access to a server farm located in your core network that is underperforming. You can be alerted to this issue so you can take action. Another futuristic model may show us being able to reroute traffic to a different server pool that is underutilized so that the load is more evenly dispersed.

Another platform type for AI is in the information technology operations management (ITOM)/information technology service management (ITSM) space. A leader in this space is ServiceNow (`servicenow.com`), another platform I will cover in depth in this book. ServiceNow is a service desk, service management, and, with ITOM, operations management platform that allows for incident triage, ticketing, asset management, and configuration management with configuration items (CIs), and ultimately it allows for AIOps through its platform. This is a more all-encompassing service management solution that performs different functions than a tool like Splunk, although Splunk may be a trusted source of data input into ServiceNow for ITOM.

Application performance monitoring (APM) is yet another platform to consider where AI is concerned. It is important to monitor and manage application performance. Although Splunk and ServiceNow offer similar functionality, an APM platform like Riverbed (`riverbed.com`) can be used to examine the end-user experience with the applications you deploy as services to them. APM will look at the services and how they perform based on AI via big data analysis it collects to make smart or automated decisions to help immediately solve customer problems.

So, as you can see, there is a myriad of offerings in different spaces. Enterprise monitoring, ITOM, and APM are the most sought-after solutions in healthcare IT and AIOps. However, other important platform choices exist, such as the ability to converge. So, as I noted before with ServiceNow, you may use this as your main ITSM tool with Splunk and Riverbed offerings added to complement infrastructure, applications, and other systems to create a complete package. Figure 1.4 shows how all of these systems (platforms) can be used together to offer one complete AIOps package in your enterprise.

Customer Experience and AIOps

All of technology is in place to make life easier, not the other way around. I know it seems to be the opposite sometimes, so we need to really focus on supplying solutions that benefit our customers in the end. When something is deployed into production, we should ask these questions: How did this help the company? Whom did it benefit in the organization? Did I get my return on investment (ROI)? Ultimately, how did this impact the customer? Ideally, the answers are positive ones.

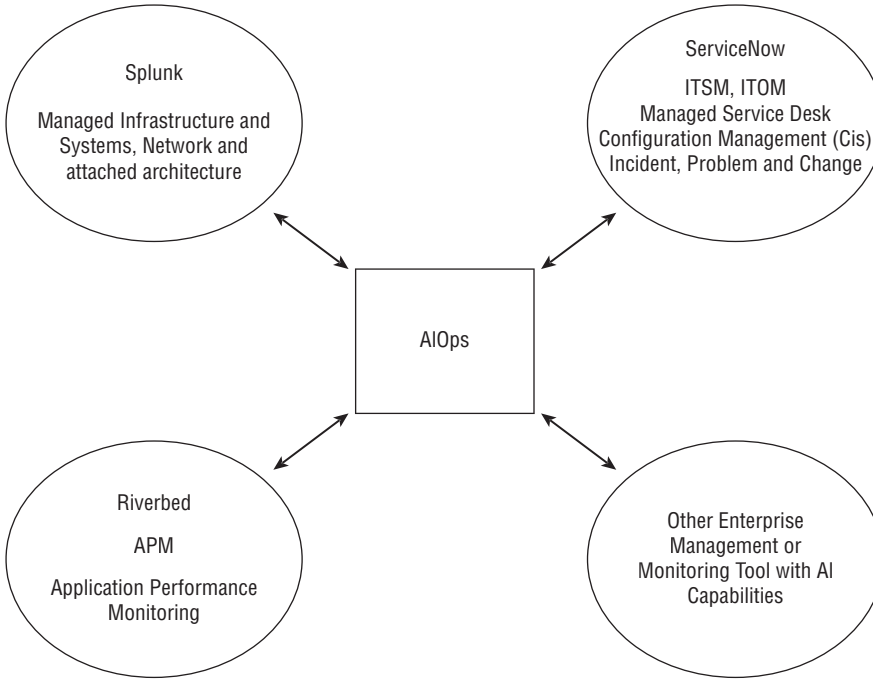

Figure 1.4: Convergence of platforms

Customer experience is the ultimate goal for all service offerings and solutions. All service delivery should be focused toward improving customer experience, and AIOps can help deliver a positive customer experience if done correctly. "Done correctly" is the key factor here. If you do not deploy these systems correctly, you never really get to the priceless benefits of full automation and decision making, and you therefore reduce your AIOps system to nothing more than a widely deployed and costly paperweight.

GETTING THE VALUE OF AIOps ON YOUR SYSTEM

Because AIOps is a pervasive overlay on your infrastructure, it impacts the performance of everything. When deciding how you want to use your AIOps platform, it is important to understand agent versus agentless technology. If you deploy an agent (an application or software program) to all systems you want to include in your AIOps deployment, it has to be considered as a possible impact to that system. We do not want to pay to deploy an AI system

that does nothing more than just monitor for issues. We hope to deploy an AIOps platform that we can take full advantage of, which includes the learning, intelligence, and automated decision-making functions that make it so valuable.

AIOps Considerations and Goals

One of the biggest goals of AIOps is customer experience, but what does that translate to? For one, it means that when you use an AIOps solution, you will be able to improve your event correlation, incident triage, speed to root-cause analysis (RCA) with preventative and corrective actions, and automated decisions that can speed mean time to resolve (MTTR), which is a key metric when considering KPIs and CSFs. Key performance indicators (KPIs) and critical success factors (CSFs) are the holy grail of metrics to executives looking to ensure the ROI of deploying costly systems like AIOps. By focusing on the KPIs, you are well on your way to proving your AIOP's ROI. Figure 1.5 shows how these metrics are used to analyze and provide for better customer service.

Key Performance Indicator (KPI)	Explanation
Number of Repeated Incidents	Number of repeated incidents, with known resolution methods
Number of Incidents	Number of incidents registered by the service desk and grouped into categories
Average Initial Response Time	Average time taken between the time a user reports an Incident and the time that the service desk responds to that incident
Incident Resolution Time	Average time for resolving an incident and grouped into categories
Resolution within SLA (Service Level Agreement)	Rate of incidents resolved during solution times agreed in SLA

Figure 1.5: Analyzing KPIs for AIOps

Your second major goal should be transformation. By improving this triage, collection, handling, and resolving of problems (especially with automation), you are fundamentally transforming your workforce, technology environment, and overall footprint in the organization. Technology should always be recognized as a service. If it is not providing value, it is a detriment to the business and its ability to thrive. Technology needs to work for the business, not the other way around. By transforming your operations to be self-aware and actionable, you are raising the bar on transforming into the next age of technology offerings, operations, and service delivery.

Next should be innovation. If you are able to focus more on self-enlightenment of your technology, you will free time cycles to build on innovative solutions. As well, if your systems are fed with real-time data, you are able to make smart decisions on that data that should also help drive innovation. All of this in turn helps to create a better customer experience. Our customers include patients, clients, clinicians, business partners, and all the people we work and interact with, but they extend beyond that. Our customers include everyone who touches technology, which in this day and age is everyone.

Summary

In this chapter, we covered the fundamental knowledge needed to understand healthcare technology and operations as well as the fusion of AI to create AIOps. We covered a lot of important foundational information required to make smart decisions based on what tools you want to get, how to deploy them with a good strategy, and why they should be implemented in the first place. The focus of this introductory chapter was to begin your journey into the much larger word of AI and ML. AI and ML become increasingly important as the technologies around them have matured. We can now use AI and ML to help us reduce outages to our critical systems, improve interconnectivity, and perform other essential functions with growing confidence. This is where the importance of AIOps begins to emerge. I also touched on the importance of the EMR and its underlying infrastructure, how interconnected these systems really are, and how any impact to them could cause negative outcomes.

This chapter discussed the healthcare market and how technology has evolved, particularly within the AIOps space. AIOps is being used to create efficiency, reduce downtime, improve response time, reduce waste, and ultimately create a better experience for all patients, clinicians, and healthcare professionals. In the next chapter, we move beyond the foundational knowledge and dig deeper into AIOps and how to use it.