cloudticity

Best Practices for Telemedicine Providers

AWS and Cloudticity combine to improve reliability, security, performance, and cost savings

Almost overnight, telemedicine firms find themselves on the front lines of the COVID-19 pandemic. As more and more patients and doctors switched from in-person visits to remote consultations, providers saw their traffic volumes soar to unprecedented levels in just days. Conventional wisdom says that's a good problem to have, but there's significant business risk in growing faster than your infrastructure can adapt. In particular, dropped calls and poor video and audio quality can lead to a poor user experience, customer churn, and even liability lawsuits.

As a leading managed services and managed security services provider specializing in healthcare, Cloudticity works with telemedicine and other providers to address reliability, security, compliance and cost optimization challenges. In the process we have identified best practices that can benefit any telemedicine enterprise — and many other healthcare organizations as well. This ebook presents a summary of our findings.

Contents

Overnight Success, Major Challenges

When the COVID-19 pandemic descended on the United States with full force, one telemedicine provider saw its video call volume increase by two orders of magnitude almost overnight. As a result, an infrastructure that worked well for a fairly stable level of baseline traffic was suddenly plagued by reliability, performance and scalability problems. Customers who had used the system reliably for a long time suddenly had a new experience and were unable to connect consistently after COVID due to insufficient performance.

Here are some of the structural obstacles that this telemedicine provider faced:

Performance

The provider had set up points of presence (POPs) around the United States to connect locally to its customers. Each POP included provider-owned hardware and software hosted in a colocation facility. When the incoming call volume spiked, the provider had no ability to react quickly—the capacity was fixed by the hardware.

Scalability

The provider used AWS resources for application hosting, database, and storage. Allocations for EC2 instances were static based on the provider's baseline traffic over the previous months. As demand spiked, the provider's IT staff manually boosted resource allocations to accommodate the increase, but had to use guesswork and continually adjust the allocations upward—no one could have predicted the magnitude of COVID-19.

Security

Over the years, developers, test engineers and others had established their own "rogue" VPNs to access information within the AWS infrastructure. The result was hundreds of individual connections that made it difficult to prevent data exfiltration—after all, a cybercriminal only had to compromise a single connection to gain access to PHI and the organization's intellectual property.

The following sections present best practices in four categories: reliability, compliance, security, and cost optimization. While these best practices are targeted primarily at telemedicine companies, they can also be applied to a wide range of healthcare providers, indeed, any organization that requires the ability to deliver vital services in the face of dramatic swings in demand. Following best practices, we offer suggestions for structural improvements that can help telemedicine providers modernize their legacy architectures.

Pandemic Perspectives

COVID-19 has pushed telemedicine into the mainstream, but will it stay there after the pandemic? The American Medical Association weighs in on this important question.

Reliability Best Practices

It's hard to think of a healthcare service that doesn't require high reliability. When lives are at stake, the need for reliable service delivery is obvious. But even ancillary services such as patient record retrieval and appointment scheduling affect the overall quality of care and therefore must be continually available — downtime for any healthcare-related service is simply not an option. This section presents tested and proven best practices to help telemedicine providers accommodate even wild swings like those produced by COVID-19:

- Enable auto scaling
- Fortify critical databases
- Offload bandwidth

Enable Auto Scaling

Best Practice

AWS Auto Scaling monitors applications using Amazon CloudWatch and automatically adjusts capacity to maintain steady, predictable performance at the lowest possible cost. For telemedicine providers, Auto Scaling adds Amazon EC2 and other resources as needed to ensure that spikes in high-bandwidth video traffic will not overwhelm capacity. When properly configured, Auto Scaling also contributes to cost optimization, which is discussed later in this document.

Pro Tip

Using scaling policies, you can exercise more precise control over the dynamic scaling process, for example, set maximum scaling limits or base scaling decisions on multiple measurements of loading. In most cases, setting up dynamic scaling requires deep knowledge of Amazon CloudWatch and policy-based process control and is not recommended for DIYers.

Services

AWS Auto Scaling

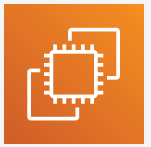

Amazon EC2

Amazon CloudWatch

Fortify Critical Databases

Best Practice

As demand soars, information queries can become a bottleneck and cause application slowdowns or outages. AWS Relational Database Service (RDS) allows you to scale your database's compute and storage resources with only a few mouse clicks or an API call.

However, scaling the database alone may not be enough, because the root problem is that too many read operations overwhelm the query capabilities of a single database instance. The answer is to launch additional read replicas of the database to spread the read load and provide consistent query response. AWS RDS has built-in capabilities for read replicas.

Pro Tip

For cloud applications, Cloudticity typically uses AWS Elasticache to create and manage the read replicas mentioned above. Amazon ElastiCache allows you to seamlessly set up, run, and scale in-memory data stores in the cloud. AWS Elasticache supports two open source code bases, Redis and Memcached.

Services

Amazon RDS

AWS ElastiCache for Redis

AWS ElastiCache for Memcached

Offload Legacy Bandwidth

Best Practice

Established telemedicine providers often rely on legacy applications that were not intended to scale rapidly enough to respond to an event such as COVID-19.

As a case in point, the provider discussed earlier had a fixed-capacity telemedicine calling infrastructure in several colocation facilities. In the short term, migrating that capability to AWS was out of the question. However, Cloudticity and the customer identified a significant subset of noncritical tasks (for example, downloading software updates) that could be safely moved to AWS using Amazon CloudFront content delivery network, reducing the load on the proprietary front end.

Pro Tip

This best practice may require changes to your existing applications. Leveraging health-care-specific cloud expertise, Cloudticity Professional Services can help you optimize your apps for the cloud and improve performance and availability while maintaining and improving HIPAA and HITRUST compliance.

When migrating customer-facing applications to AWS, telemedicine providers may need to maintain existing IP addresses for some of their customers. Bring Your Own IP (BYOIP), a little-known feature of Amazon EC2, enables you to use an existing IPv4 address for the AWS-hosted application.

Services

Amazon CloudFront

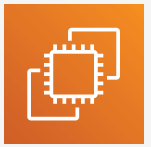

Amazon EC2
(BYOIP feature)

cloudticity

Professional Services

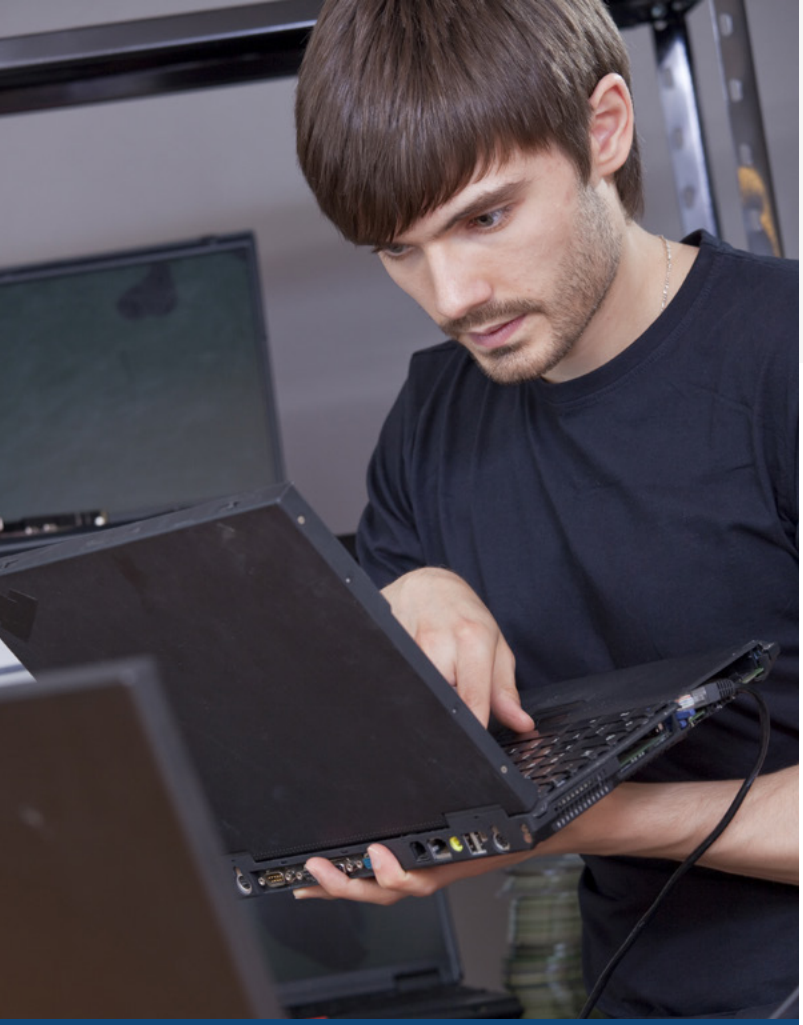

Pandemic Perspectives

Whenever patient information is involved, HIPAA comes into play. The HIPAA Journal offers HIPAA compliance guidelines for healthcare providers as they adapt their business processes to the unique and challenging world of telemedicine during COVID-19.

Compliance Best Practices

Cloudticity views compliance not as an annual checklist — an annoying task to be passed and forgotten until next year — but rather as a real-time process that drives security. We focus exclusively on the healthcare market and as a result have significant expertise with all healthcare regulations including HIPAA, HITRUST, and more. Our recommended compliance best practices include:

- Automate compliance remediation
- Perform formal risk assessments
- Audit network attack surface

Automate Compliance Remediation

Best Practice

Healthcare applications cannot offer high reliability if problems are remediated at human response speed — the numbers just don't add up. The need is to automate everything that can be automated and rely on people to handle exceptions that don't impact reliability. The Cloudticity Oxygen platform uses AI and operational intelligence to automatically remediate anomalies when possible, and presents clear choices to operations staff for cases that cannot be resolved automatically. This approach shortens time to resolution, avoids downtime and improves reliability of telemedicine service delivery. Automated response also plays a critical role in security, as discussed below.

Pro Tip

Automation can only go so far — there are always interventions that require a human touch. Oxygen typically can automate about 98% of remediation tasks. For the rest, it provides a recommended action and offers the human operator three choices:

- Do Nothing: Take no action, delete ticket after 7 days.
- Fix Now: Implement the recommended actions immediately.
- Schedule: Perform the recommended actions during the next maintenance window.

Read this Cloudticity blog to see how Oxygen uses AI and automation to power intelligent service delivery.

Services

Cloudticity Oxygen

Perform Formal Risk Assessments

Best Practice

Understanding your risk level and identifying specific risk issues are critical components for developing an effective compliance plan. In our experience, many providers of healthcare services underestimate their level of risk, in part because risk is difficult to quantify.

As a good first step, Cloudticity recommends using Amazon Inspector, an automated security assessment service that helps improve the security and compliance of applications deployed on AWS. AWS Inspector automatically assesses applications for exposure, vulnerabilities, and deviations from best practices. After performing the assessment, AWS Inspector produces a detailed list of security findings prioritized by level of severity.

Pro Tip

AWS Inspector determines whether your ports are reachable from the internet through an internet gateway (including instances behind load balancers), VPC peering connections, or VPNs through a virtual gateway. These findings also highlight network configurations that allow for potentially malicious access, such as mismanaged security groups. Cloudticity Professional Services can help you interpret the results of AWS Inspector and identify specific actions needed by priority.

Services

Amazon Inspector

cloudticity

Professional Services

Reduce Network Attack Surface

Best Practice

To provide secure access to sensitive information, hybrid architectures need a VPN gateway between on-premises and AWS cloud resources. However, developers, test engineers, remote employees and others who need access to cloud-based PHI may bypass the VPN gateway by either 1) “punching a hole” in the cloud firewall to allow direct unencrypted internet traffic or 2) using VPC peering connections, which are opaque to IT and compliance groups (see diagram).

Amazon WorkSpaces, a managed and secure desktop-as-a-service (DaaS) solution, provides an elegant way to allow access to PHI on AWS without exposing the connections to cyberattackers.

Pro Tip

Amazon WorkSpaces is deployed within an Amazon Virtual Private Cloud (VPC), providing each user with access to persistent, encrypted storage volumes in the AWS Cloud using AWS Key Management Service (KMS). You can provision Amazon WorkSpaces to include all the tools your developers need to build applications quickly.

Read this Cloudticity blog for more information about using Amazon WorkSpaces.

Services

Amazon Workspaces

Amazon Virtual Private Cloud

Amazon Key Management Service

Pandemic Perspectives

As telemedicine traffic explodes in response to lockdowns, so does the volume of hacking attempts. The consequences of a breach are chilling: reputation loss, customer defections, regulatory fines, lawsuits, and more. HealthIT Security outlines key must-have privacy and security needs during the pandemic.
Many telemedicine applications rely on VPNs. The US Department of Homeland Security has published guidelines for VPN security during the COVID-19 pandemic.

Security Best Practices

Telemedicine providers are understandably worried about how cyberattacks can cause havoc in their businesses. Imagine a malware-caused outage--real or threatened--during remote robotic eye surgery or hackers eavesdropping on intimate doctor-patient conversations. While the full scope of security best practices is beyond the scope of this paper, here are three ways that telemedicine providers can bolster their security profile:

- Deploy proactive security
- Encrypt data storage
- Harden operating systems

Deploy Proactive Security

Best Practice

Modern cyberthreats are becoming steadily more sophisticated in evading traditional security measures and more devastating once they penetrate the network perimeter. For that reason, telemedicine providers need a highly proactive approach to prevent malware-based outages, theft of intellectual property, and exfiltration of personal health information (PHI).

To meet this need, Cloudticity partners with Trend Micro, a leader in the security space. Specifically, Trend Micro Deep Security integrates key security features into the Oxygen platform including anti-malware, application control, and IPS. As a result, Oxygen can detect suspicious system changes in real time, isolate and quarantine the resource, and prevent the spread of exploits by locking down any server whose configuration differs from the installed settings.

Pro Tip

No single security measure can adequately secure your telemedicine infrastructure — you need multiple layers. In addition to proactive security, Cloudticity recommends storage volume encryption and operating system hardening, which are described in the next two sections.

Services

Trend Micro
Deep Security

Cloudticity Oxygen

Encrypt Data Storage

Best Practice

Data encryption is the last line of defense for PHI and other critical information. Even if a hacker can elude perimeter and proactive network security and exfiltrate data from the provider, that data is useless to the hacker if it is encrypted.

For that reason, Cloudticity recommends Amazon EBS encryption for all web and application servers running on Amazon EC2 instances. EBS encryption is available for most modern EC2 instance types. Check the online list of supported instance types to learn more.

Amazon EBS encryption uses customer master keys from AWS Key Management Service when creating encrypted volumes and snapshots. Encryption operations occur on the servers that host Amazon EC2 instances, ensuring the security of both data-at-rest and data-in-transit between an instance and its attached EBS storage.

Pro Tip

For additional protection, you can encrypt database instances and snapshots at rest by enabling the encryption option that comes standard with Amazon RDS. Data that is encrypted at rest includes the underlying storage for DB instances, its automated backups, read replicas, and snapshots.

Services

Amazon EBS
(encryption)

Amazon Key
Management Service

Amazon RDS
(encryption)

Harden Operating Systems

Best Practice

Microsoft Windows Server and Linux are attractive targets for cybercriminals because they provide complex capabilities, frequently remediate vulnerabilities, and are ubiquitous (providing a good chance of finding an unpatched system).

Cloudticity Oxygen defends against OS attacks by providing their customers hardened images of Windows Server and Linux. These images are highly secure by default using configuration best practices provided by the Center for Internet Security. Gaining administrative access becomes extremely difficult, and even if hackers are able to succeed in altering OS settings or file system data, Oxygen immediately detects the deviation and quarantines the server as described earlier in Deploy Proactive Security.

Pro Tip

Hackers use OS-based techniques such as remote code execution and elevation of privilege to take advantage of OS vulnerabilities for their nefarious purposes.

Cloudticity's hardened OS images not only assist with security but also streamline the process of provisioning new OS instances.

Services

CLOUDTICITY OXYGEN™

Cloudticity Oxygen

CIS.

Center for Internet Security

Pandemic Perspectives

COVID-19 puts tremendous pressure on the US medical system, which was already struggling to manage the cost of healthcare services. In response, Medicare, Blue Cross, and other healthcare payers have added coverage for telemedicine costs that were previously not covered.

Cost Optimization Best Practices

As many providers have learned the hard way, costs for AWS services can quickly mount up to budget-breaking amounts. Fortunately, with the right tools, experience and expertise, it is possible to optimize costs without impacting performance, as this section shows.

- Rightsize AWS instances
- Use reserved instances
- Leverage AWS savings plans
- Optimize container spend

Rightsize AWS Instances

Best Practice

Right sizing application instances is the most effective way to control cloud costs. It involves continually analyzing instance performance and usage needs and patterns and then turning off idle instances and right sizing instances that are either overprovisioned or poorly matched to the workload. Use these AWS tools to streamline monitoring and analyzing historical data:

- Amazon CloudWatch allows you to observe CPU utilization, network throughput, and disk I/O and match observed peak metrics to a new and cheaper instance type.
- AWS Compute Optimizer recommends optimal AWS Compute resources for your workloads to reduce costs and improve performance by using machine learning to analyze historical utilization metrics.
- AWS Trusted Advisor lets you inspect your AWS environment to identify idle and underutilized resources and provides real-time insight into service usage.

Services

Amazon CloudWatch

AWS Compute Optimizer

AWS Trusted Advisor

Use Reserved Instances

Best Practice

For use cases in which workloads and resource needs are more predictable, you can cut costs significantly — up to 72% — by using Amazon EC2 Reserved Instances (RI) instead of On-Demand resourcing. However, RIs are a contractual commitment and cannot be revoked.

Three variables drive the AWS pricing algorithm: instance attributes, term commitment, and payment options. Reserved Instance Reporting, a feature of AWS Cost Explorer, provides reports that help you manage your organization's RI spending.

Pro Tip

If your circumstances change and you no longer need your reserved capacity, consider selling your RIs to a third party via the Amazon EC2 Reserved Instance Marketplace. Some restrictions apply, contact Cloudticity Professional Services for more information.

Read this Cloudticity blog to better understand best practices for buying AWS RIs.

Services

AWS Cost Explorer

Reserved Instance Reporting

cloudticity

Professional Services

Leverate AWS Savings Plans

Best Practice

AWS has introduced two new AWS Savings Plans this year that offer substantial savings compared to On-Demand compute pricing. Both plans require either a one-year or three-year commitment to a minimum spend per hour and can be used in conjunction with Reserved Instances (RIs).

EC2 Instance Savings Plans are well suited for companies with predictable workloads and larger installations of specific instance families. These plans offer the same savings as RIs and are purchased for a specific region and family, for example, t3 or m5.

Compute Savings Plans provide slightly less savings than EC2 plans but provide more flexibility — they can apply to EC2 Compute and Fargate usage regardless of region or instance family.

Pro Tip

Because AWS offerings and terms are dynamic, setting up your infrastructure to take full advantage of savings in the AWS pricing structure requires extensive recent experience and expertise with AWS cost optimization.

Read this Cloudticity blog to understand the differences between the Savings Plans as well as the pros and cons of each.

Services

AWS Savings Plans

Professional Services

Optimize Container Spend

Best Practice

You can optimize your spend for containerized applications using AWS Fargate, a serverless compute engine that provides on-demand, right-sized compute capacity for containers running in clusters managed by Amazon Elastic Container Service (ECS) or Amazon Elastic Kubernetes Service (EKS). With Fargate, you only pay for the resources required to run your containers, eliminating overprovisioning in container environments.

Pro Tip

Implementing Dockers or Kubernetes clusters requires resources and expertise that most companies lack. Even seasoned Kubernetes expert can take hours to set up a basic Kubernetes cluster manually. Add in the need to implement proper security, alerting, health checks, and dashboards, Kubernetes management quickly becomes a full-time job.

Cloudticity Managed Kubernetes, part of the Cloudticity Oxygen platform, offers a cost-effective alternative. With Cloudticity, your team can quickly and easily create and configure clusters without deep knowledge of Kubernetes. Based on EKS, Cloudticity Managed Kubernetes provides options for creating clusters including hardened images, instance types selection, a choice of managed node groups or Fargate profiles, and public or private API endpoints--all in a single deployment workflow.

Services

AWS Fargate

Cloudticity Oxygen

Pandemic Perspectives

Given the urgency of the COVID-19 pandemic, healthcare providers need to extract insights from information faster and more accurately than ever before. Some have adopted the Cloudticity Healthcare DataHub to improve COVID-19 surveillance and inform the response and recovery.

Structural Improvements

While the following recommendations do not qualify as best practices, they represent strategic initiatives that telemedicine providers have adopted to reduce operating costs, drive revenue growth, and gain competitive advantage:

- Modernize monolithic applications
- Migrate call infrastructure to AWS
- Implement Infrastructure as Code

Modernize Monolithic Applications

Although recent events have shone a spotlight on the technology, telemedicine has been in use for decades. As a result, established telemedicine providers often have significant investment in business-critical monolithic applications that are difficult to adapt to the cloud environment. Simply moving these code bases to the cloud — often called "lift and shift" — offers only marginal benefits. To gain the full range of cloud benefits, applications must be rearchitected from monolithic to microservices.

Completely rewriting a monolithic application all at once is usually out of the question due to the costs involved and the lack of staff with the required skill sets. A more palatable approach is a stepwise migration to a microservices and containerized architecture. Using this methodology, the development team rewrites key modules as microservices and then modifies the monolithic application to point to these cloud-native functions.

One advantage of this approach is that organizations can transition at their own pace without disrupting ongoing business. Cloudticity has extensive experience helping customers transition from monolithic applications to microservices using AWS services such as those shown below.

Selected AWS services used in microservice architectures:

Amazon Elastic Container Service

Amazon Elastic Container Registry

Amazon Elastic Kubernetes Service

Amazon VPC

Elastic Load Balancing

AWS Cloud-Formation

Amazon CloudWatch

AWS Lambda

cloudticity

Contact Cloudticity Professional Services to schedule an evaluation of your applications.

Migrate Call Infrastructure to AWS

Every telemedicine provider has a front end that manages the telecommunications connections between healthcare provider and patient. When the COVID-19 pandemic caused call volumes to spike overnight, many telemedicine organizations discovered to their dismay that their current calling infrastructures could not scale to meet the need.

Cloudticity helps telemedicine providers transform their rigid calling infrastructures into highly scalable cloud-native infrastructure using Amazon Kinesis Video Streams and Web Real-Time Consortium (RTC), an open-source project managed by the World Wide Web Consortium (W3C) that enables audio and video communication within web pages.

This combination allows providers to securely live stream media or perform two-way audio or video interaction between any camera IoT device and WebRTC-compliant mobile or web players.

Services

Amazon Kinesis Video Streams

Amazon Connect

World Wide Web Consortium

Implement Infrastructure as Code

Infrastructure as Code (IaC) is a methodology for managing infrastructure in a descriptive model, using the same versioning as DevOps team uses for source code. Like the principle that the same source code generates the same binary, an IaC model generates the same environment every time it is applied. Learn more about IaC here.

IaC is particularly important for AWS, a platform that is constantly evolving using DevOps practices such as continuous integration and continuous delivery (CI/CD). A key benefit of IaC is that it captures and replicates provisioning settings without relying on default settings, which can change over time. In a security context, IaC allows security teams to find unwanted changes in application configuration and quickly revert to a known secure state.

Cloudticity implements IaC using AWS CloudFormation. CloudFormation provides a common language for modeling and provisioning AWS and third-party application resources in cloud environments. With CloudFormation, you can use programming languages such as YAML or a simple text file to model and provision all the resources needed for your applications across all regions and accounts.

Services

AWS CloudFormation

Oxygen on AWS

Cloudticity is not only an AWS Partner, we are an AWS customer as well. Our Oxygen platform incorporates multiple AWS services including AWS Lambda, Amazon Kinesis, Amazon Athena, Amazon Elastic Map Reduce (Amazon EMR), Amazon QuickSight, and Amazon EC2 Systems Manager.

AWS Services Incorporated into Cloudticity Oxygen Platform:


AWS Lambda


Amazon Kinesis


Amazon Elastic Kubernetes Service


Amazon Elastic Map Reduce


Amazon QuickSight


Amazon EC2 Systems Manager


AWS Config


AWS SNS


Amazon CloudWatch


Amazon S3


Amazon Dynamo DB

Cloudticity chose AWS for the same reason that we recommend AWS to our customers: their services are scalable, flexible, secure, and cost-effective.

That decision helps us streamline the Oxygen development cycle and allows us more time to do what we do best: Help our customers build and maintain reliable, secure, compliant, scalable, and cost-effective infrastructures.

Why Cloudticity?

Cloudticity provides managed offerings for IT services, security and compliance as well as consulting services for cloud migration, application optimization, DevOps automation, and HITRUST CSF Certification. What sets us apart from other managed services providers is our emphasis on automation and AIOps instead of human intervention. To that end, we developed Cloudticity Oxygen™, our proprietary managed services delivery platform that leverages automation and artificial intelligence to oversee all aspects of operations for healthcare systems and applications. Our recommended best practices include Oxygen and Cloudticty Professional Services where relevant and we do so unabashedly — both have proved themselves in a wide range of healthcare use cases.

Our entire practice focuses on healthcare. Cloudticity is an AWS Advanced Consulting Partner with certified competencies in healthcare and DevOps. Our architects and consultants are experts in the use of AWS resources and how to integrate them into existing infrastructures for maximum benefit. All of the best practices identify the AWS services that are critical aspects of the recommendations.

AWS services are relatively easy to procure through the self-service portal, however, optimizing a complex architecture consisting of multiple AWS services and on-premises infrastructure requires a level of expertise and experience that most healthcare organizations lack. By partnering with Cloudticity, you get the benefit of our decades of AWS and infrastructure experience — what you need, when and where you need it.

For more information about using AWS for telemedicine applications, schedule a free consultation with a telemedicine cloud expert.

Cloudticity is reinventing the MSP operating model with AIOps and other leading-edge technologies. To learn more, read the Cloudticity white paper **AIOps and Automation: Keys to Intelligent Managed Service Delivery**.


cloudticity
Ready to future-proof your telemedicine application?
LEARN MORE