Information Security White Paper

Introduction to Medidata's Information Security Program

A mature and effective Information Security program is critically important to the integrity of clinical drug trials. Medidata has developed and implemented a consistent and validated enterprise program based on the widely recognized NIST 800-53 security framework. The integrity of this framework is evaluated and confirmed by multiple annual independent external assessments, including NIST 800-53v5/FISMA audits, SOC 1 Type II and SOC 2 Type II (with additional controls) audits, as well as annual ISO 27001:2013, 27018:2014 and 27701:2019 assessments.

Our information security program is the foundation of our industry-leading unified, intelligent platform that supports your end-to-end clinical development. Medidata's secure clinical cloud computing solutions deliver high availability, integrity, confidentiality, reliability and the flexibility to enable our wide range of applications.

With the acquisition by Dassault Systèmes, Medidata is perfectly positioned to take advantage of the global scale that Dassault provides while continuing to deliver secure, stable and scalable infrastructure.

This White Paper provides an introduction to understanding the investment and expertise Medidata brings to the forefront, in order to protect the data entrusted to us.

Trust and Transparency: The Medidata Trust Center

Medidata's Information Security program, along with our Data Privacy and Global Compliance functions, ensure that Medidata has created a secure, stable, and scalable cloud platform, robust data governance processes, and an inspection-ready Quality Management System — which, when you put them all together, are critical enablers to success in clinical trial execution.

We are proud of what we do and are happy to show it. So, in addition to this document, we post quarterly vulnerability summaries, penetration tests results, certifications, audits, and other security related matters to our Trust Center at https://www.medidata.com/trust so that our customers and their clinical sites can have a level of comfort in that the protections are what the patient expects.

For any questions or clarifications, feel free to reach out to asksecurity@medidata.com.

"Information Security is crucial to the protection of patient privacy and in some cases their very life; when I became a clinical trial patient, the protection of my healthcare data became personal. I can tell you that as a trial patient and security professional who understands the cyber threats, Medidata is the only firm I want storing and processing my information.
"

- Glenn Watt, Medidata Chief Information Security Officer (2007-2018)

Program Overview

We at Medidata believe that Information Security is a program, not a checklist. Using a "Security-by-Design" and "Privacy-by-Design" philosophy, integrity in our systems is built in at the ground floor.

Medidata has a team of dedicated Information Security professionals, with over 150 years of accumulated Information Security experience in practicing InfoSec in the Life Science, Technology and Defense industries.

Reporting to the Chief Technology Officer / Chief Information Officer / Head of Product, our InfoSec team is responsible for the full and independent management of the InfoSec program, which is based on the CoBIT, ISO, and SOC standards. This program is audited at least four times a year by independent, third-party auditors including PricewaterHouseCoopers, Apex CyberTek, and DQS.

Infrastructure vulnerability scans, peer code reviews, static source code analysis, dynamic scanning of URLs are all performed at least monthly or during any significant change in the environment. Additionally, prior to any product being released for General Availability, every product is scanned for vulnerabilities. We also run an annual Red / Blue / Purple team engineering test which has the penetration testers working alongside the response team, in order to maximize learnings.

The core of the environment is hosted within Amazon Web Services (AWS) data centers, in the US-East, US-West, Frankfurt, Ireland and Paris for truly global coverage using a homogenous support and security model. We use all of Amazon's security features including tight security groups, rigid network access control lists, CloudFront, AWS Shield and Advanced Shield, CloudTrail, as well as the organic security built into all Amazon's products.

We also have our private cloud-based operation in Houston, Texas with an alternate processing site in Frankfurt, Germany in order to protect the systems in the event of a regional disaster.

In order to provide 24x7x365 support, we use a Managed Security Service Provider that collaborates closely with our security operations center (SOC) to provide protection against any and all threats, regardless of the time of day.

We are proud of the trust that our customers and their patients have provided to us, and we earn it every single day.

Cloud vs On-Premise Security

Security in the cloud is not managed the way it is in a legacy premise-based environment. The concept of bastion perimeter is more nebulous; without a single network connection to the outside world, you have to harden everything.

Cloud Security has matured beyond traditional data center security; rather than relying on a set of controls at the perimeter, controls and technology permeate the entire infrastructure and application layers. Where you had a perimeter firewall, you now have security groups, Network Access Control Lists (NACL), File Integrity Monitoring, Malware Protections, Log Inspection, Web Reputation checking and Web Application firewall layers on each and every host.

The hypervisor segregation (see: Figure 1) allows for completely independent hosts and databases but converged to allow to take advantage of the advanced security features such as anti-beaconing monitoring, cloud access security brokers and web application firewalls.

Figure 1

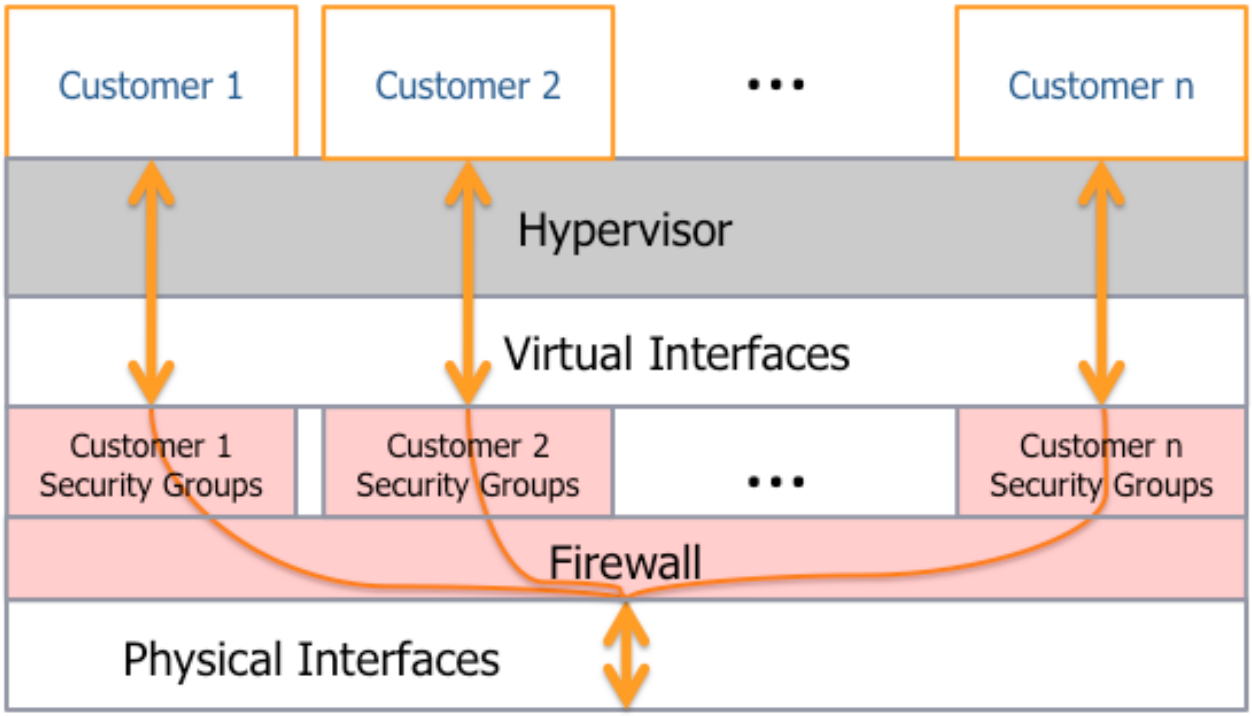

Figure 2

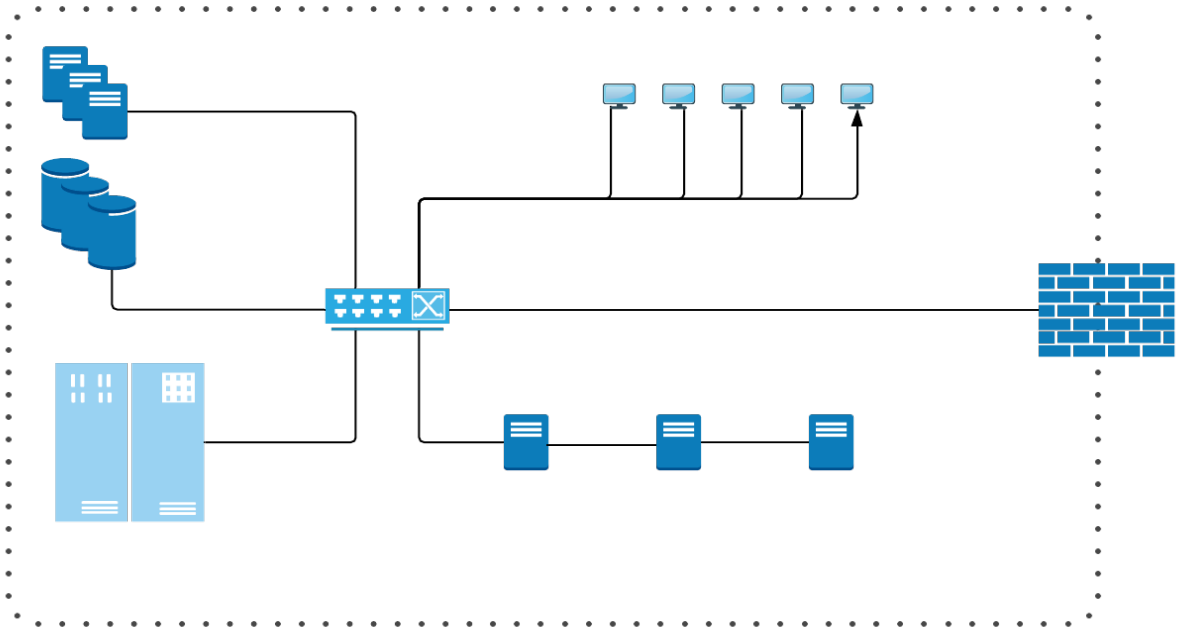

In a legacy premise-based environment (see: Figure 2), security services are enabled at the network perimeter, and where technology exists, installed on the hosts. The challenge is that even with management software, consistency and speed of deployment is problematic. Once a malefactor is inside the perimeter, it is hard to track and stop them from gaining more access.

Figure 3

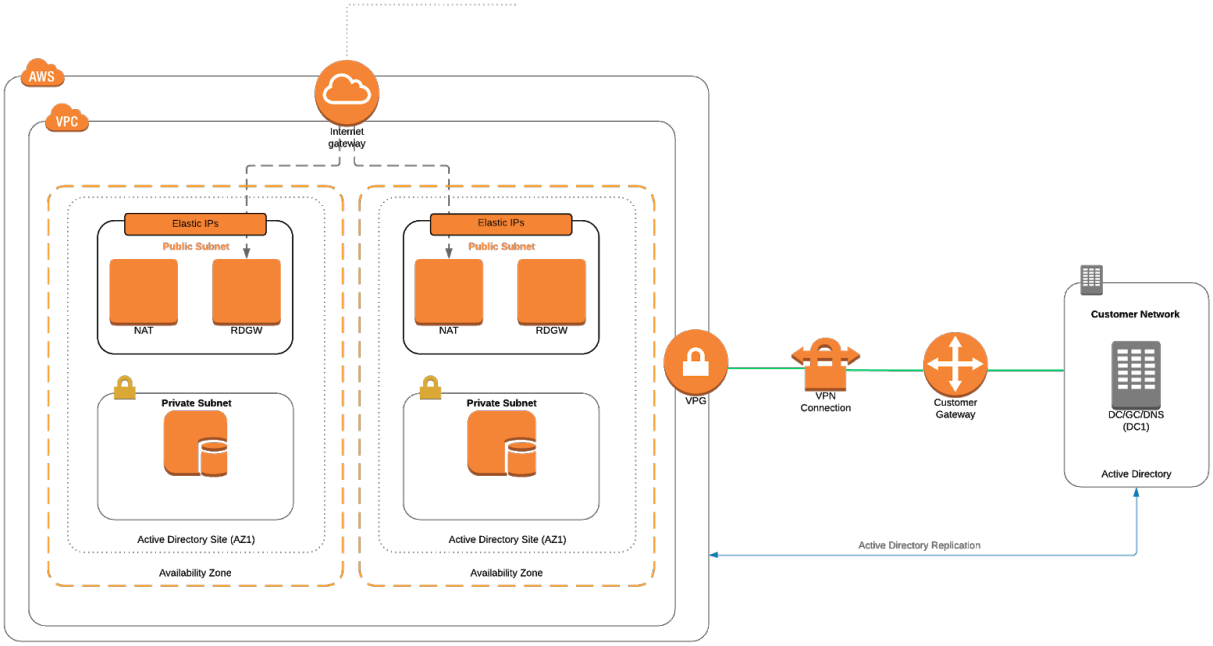

In a cloud environment (see: Figure 3), security permeates the fabric of the infrastructure; with security control at every single layer, consistently applied, and centrally managed. This allows bringing new services to bear simultaneously across the entire environment, quickly and accurately. With detailed instrumentation, a malicious actor can be caught and contained, before there is a security event. Updates to any layer can be done in a matter of minutes and is much easier to manage and report against.

Tenancy

Tenancy refers to software application architecture describing the relationship between the customer *(tenant)* accessing the software and the number and distribution of instances of that application. In a **Multiple Instance Single Tenant ("MIST")**, each instance of an application is only accessed by one tenant. Therefore, each client accesses a unique application instance with a logically partitioned database. Segregation is achieved using **Attribute Based Access Control** (**ABAC)** and **Role Based Access Control (RBAC)** that provide strong separation of customer data.

In a **Single Instance Multi-Tenant ("SIMT")** environment, a single instance of an application operates and is accessed by many tenants. With SIMT applications, all customers use the same application while the customer data remains in logically separate databases. Importantly and fundamentally, SIMT applications are architected and designed such that customers cannot access data and configuration information that belongs to other customers.

SIMT applications inherently carry less risk to the customer than MIST applications primarily because of the consistency of controls applied, and speed of response for changes and controls. An upgrade or patch is applied to a single point, and applied uniformly which provides a seamless barrier to malicious activity.

Developed using the same rigorous Medidata standards as MIST applications, customers adopting SIMT products realize immediate benefit from ongoing product enhancements — including enhancements driven by regulatory changes — and feedback from Medidata's global client base. SIMT application upgrades are fast and highly effective.

Governance and Management

The best Information Security programs are consistently supported from the top, starting with the board, and permeating the entire leadership of an organization. The executives have the right background and pedigree to ask the right questions and provide proper direction with respect to the Information Security program and how it meets the organization's needs.

Medidata adopted an ISO 27xxxx style Information Security Management System (ISMS), which drives a three-year strategy. Using NIST 800-53 as the set of governing principles, guidance provided to the InfoSec team comes straight from the board of directors, maintaining independence from operational matters for maximum objectivity. The evolution of the ISMS includes the addition of the adjunct Privacy Information Management System (PIMS), which mutually support one another.

Our NIST 800-37 based Enterprise Risk Management program goes beyond InfoSec, and helps manage risk in the Human Resources, Operations, Privacy and other areas of the business.

With respect to the InfoSec team, it consists of collective experience of over one hundred fifty years in Life Sciences, Technology and Information Security. There is tenure in the National Security Agency, the United States and Royal Navy as well as other heavily secured organizations.

The InfoSec team is composed of highly credentialed individuals holding diverse certifications and skill sets including, but not limited to:

- Certified Information Systems Security Professional (CISSP)
- Certified Incident Handlers (CIH)
- Forensic examiners
- Offensive Security Certified Professional (OSCP)
- AWS Architects
- CISA / CISM / CRISC
- ISO 27xxx Lead Auditors
- Payment Card Industry Internal Security Assessors (PCI-ISA)

From top down, Medidata's Information Security team, which is part of the Dassault Systèmes team, has complete management support and guidance, which allows the team to bring to bear centuries of experience, skills and techniques to the job.

Technical Controls

Technical controls are key success criteria in protecting patient data. Medidata makes significant investments in cutting edge technologies, coupled with our "best-in-breed" security techniques and practices, provides a secure, stable and scalable architecture.

We categorize these controls into the following groups:

- Policy, Physical Security and Training
- Data Protection
- Availability
- Endpoint Security
- Network Security
- Defect and Vulnerability Management
- Identity Management
- Monitoring
- Independent Testing and Review

The holistic design provides for defense-in-depth, which allows for a control to fail, while maintaining the confidentiality, integrity and availability of the systems.

Policy, Physical Security and Training

Policy

Policy is the guide to behavior, and without it, employee, contractor and vendor responsibilities would be unclear and imprecise at best. Our policies are based on commonly accepted frameworks including CobIT and ISO 27xxx. We monitor activity closely, and enforcement is strict and consistent. We train each and every employee and contractor, at hire, and annually on the entire policy hierarchy, to ensure that the duty and responsibilities are clear and understandable.

Physical Security

All our data and systems are housed in TIA Level 3+ data centers (the highest data center security standard that Amazon has) in order to provide state-of-the-art protections at the front door. All data centers are unmarked with unpublished addresses, cameras with digital recorders, 24x7 uniformed guards, biometrics, mandatory photo-id smart cards, environmental sensors and more. Our corporate sites are similar, with tight access control uniformly throughout the entire environment. Physical security controls are described in greater detail in the Availability section below.

Education

Training is something that is central to an effective Information Security program; without it, the technical controls cannot be brought to bear to protect the data of patients and other sensitive information.

Each and every employee is trained in an all-day “New Hire'' program, which is the first exposure a new employee has to the company’s Information Security practice and is typically led by one of the InfoSec leaders. This session is intended not only to familiarize the new hire to their responsibilities, but also emphasize the role of an employee in providing protections against insider risk, ransomware, social engineering, proper use of assets, and other related items.

Developers get an extra four hours of Secure Code training, annually, which includes OWASP top 10 and SANS top 25 in order to minimize the risk of weak secure-coding practice.

If the role has access to sensitive information, or is a “control role”, additional training on their obligations and responsibilities is also held which can include high authority ID management, PII management, Data Governance policies, and more.

All of the online training is conducted via Medidata’s Learning Management System, and employees and contractors are held accountable for timely completion. This includes annual awareness courses for all staff created by KnowBe4, a leading InfoSec training platform.

We refresh the education at least annually, in order to ensure that we make the internal community aware of the most recent and relevant risks.

Data Protection

Data Protections are in many layers, from the core where the data is stored, through the entire private cloud infrastructure. These data protections start at encryption, both where the data is stored at rest and while it is transmitted. Our philosophy of "Encrypt Everywhere", which provides the most flexibility and safety, using the Advanced Encryption Standard algorithm, using 256-bit keys providing military grade protection across the entire environment.

Encryption-at-Rest

Encryption is enabled at the storage unit level and is affected through hardware. For the Rave EDC data stores, the Hitachi Storage Area Network uses 256-bit Advanced Encryption Standard (AES) keys, using a proprietary key management system. For our multi-tenant systems, we also use AES-256, but use Amazon's KMS Product.

Encryption-in-Flight

All data transmission, whether internal or external, is encrypted using Transport Layer Security (TLS) version 1.2. We are evaluating TLS version 1.3; although it is not fully implemented across the industry, Medidata is ready to move to the new protocols, as soon as they are proven to be robust, secure and fully supported.

We also manage our cipher suites used in transmission tightly. Over the next twelve months, we will look to future proof (quantum protect) our ciphers. We will be deprecating the use of cipher suites that do not support forward secrecy such as the TLS_RSA_WITH_AES suite of ciphers.

The current full list is:

- TLS_ECDHE_RSA_WITH_AES_128_CBC_SHA256 (secp256r1)
- TLS_ECDHE_RSA_WITH_AES_128_GCM_SHA256 (secp256r1)
- TLS_ECDHE_RSA_WITH_AES_256_CBC_SHA384 (secp256r1)
- TLS_ECDHE_RSA_WITH_AES_256_GCM_SHA384 (secp256r1)
- TLS_RSA_WITH_AES_128_CBC_SHA256 (RSA 2048)
- TLS_RSA_WITH_AES_128_GCM_SHA256 (RSA 2048)
- TLS_RSA_WITH_AES_256_CBC_SHA256 (RSA 2048)
- TLS_RSA_WITH_AES_256_GCM_SHA384 (RSA 2048)

Data Loss Prevention

Movement of data is monitored real-time, using Data Loss Prevention tools and techniques. We also put proprietary tools in place to monitor data as it is stored and processed. Any suspicious movement of data alerts our InfoSec team automatically, and the data movement is locked down and investigated promptly.

Any data movement is governed by our ISO 27018:2014-tested Data Governance program, directly supervised by our Chief Privacy Officer, as well as heads of Information Security and Global Compliance and Strategy to ensure that data protections remain intact.

Availability

The Medidata Rave Clinical Cloud utilizes a hybrid hosting architecture, combining Private Cloud (traditional, on-premises data centers) with secure, Public Cloud (Amazon Web Services) services.

Hosting Environment Qualification

These cloud hosting environments are Qualified at multiple levels, beginning with the evaluation of any third-party facilities or services, and including the Qualification of host infrastructure and Hosting Environment, and the Validation of Medidata developed software. Through this process, Medidata develops and maintains a Hosting Environment Qualification Plan for each environment. The objective of Hosting Environment Qualification is to define the current state of the infrastructure services used to provide Hosting Environments, monitor those services and provide a blueprint for the future development of the technological landscape.

Qualification of each Hosting Environment is dependent on the evaluation of the providers for any sub-contracted services utilized by Medidata, through Medidata's 'Supplier Evaluation' process. In addition to the Qualification of Hosting Environments, Medidata performs a separate Qualification of the hosting systems of the Medidata products.

Private Cloud Hosting

In our Private Cloud environments, Medidata contracts with a property management supplier to provide services related to the physical building structure, and essential services (such as building security, primary and secondary power, fire suppression (for common areas), cooling (for common areas), and floor space/structural integrity. Medidata assumes responsibility for all systems within our suite or cage.

Houston Data Center

Medidata's Houston Data Center (HDC) is our primary Hosting Environment for our Private Cloud Hosting operations in the United States. Medidata maintains a private data center suite in a facility managed by Netrality Properties.

Netrality Properties provides Medidata with the following services:

- Primary power feed
- Backup power generator
- Facility temperature control
- Chilled water
- Physical Security
- Floorspace (Private Suite)

Within our Data Center Suite, Medidata has implemented the following services:

- Power Stabilization and UPS
- Liebert Chillers
- Rack Based Cooling
- Biometric Access Security System (biometric identifier, PIN code unique to the individual and access card.)
- Security Cameras
- Fire Suppression (FM 200)
- Commercial grade ISP router

Netrality Properties services carry the following 3rd party attestations and certifications:

- ISO 27001:2013
- ISO 9001:2008
- SOC1 Type 2
- ANSI/TIA 942-B Certification Rating 3

Frankfurt Data Center

The Frankfurt Data Center, managed by Equinix, serves as the Disaster Recovery site for the Houston Data Center.

Equinix provides Medidata with the following services:

- **Physical Security** - Equinix continually monitors (and records) all entrances and exits to the facility. All visitors must be pre-registered with Equinix and are admitted after being identified through an intercom system.
- **Climate Control -** Equinix provides management of all environmental control systems at their facility.
- **Power** - Equinix has designed electrical power delivery systems to provide an uninterrupted supply of electrical power through various primary and secondary supply mechanisms.
- **Cross Connects** - Cross Connects permit Customers to connect their equipment to other Equinix customer equipment or Equinix interconnection exchanges located within an IBX or between IBXs on a single Equinix campus.
- **Fire Suppression** - Equinix employs a complete suite of fire-focused monitoring and data center command and control.
- **Cabling -** Equinix is responsible for all cabling, outside of the Medidata's Private Cage, and is subject to Equinix's Global IBX Policy governing Network and Telecommunications.
  All cabling within Medidata's Private Cage is the responsibility of Medidata and is subject to Medidata's Network Patch Cable Standards.
- **Facility Maintenance -** For all Data Center Equipment, Equinix maintains maintenance plans, and logs of executed maintenance, in hard-copy (in note-binders) and as scanned documents in file storage.

Within our Data Center Suite, Medidata has implemented the following services:

- **Physical Security -** Medidata's Private Cage (within the Colocation Area) is restricted to authorized personnel by means of a card reader on the Cage door, using the internal access card. Within Medidata's Private Cage, Equinix maintains a secure cabinet for Medidata, with access to the secure cabinet restricted only to authorized personnel by means of an additional locking mechanism on the cabinet door

Additionally, Equinix services carry the following 3rd party attestations and certifications:

- ISO 27001:2013
- ISO 9001:2008
- SOC1 Type 2
- TIA-942:2010 Rating 3

Public Cloud Hosting

Amazon Web Services

This section provides a brief explanation of Amazon Web Services infrastructure. Please refer to Amazon Web Services Global Infrastructure website here for complete documentation.

AWS Regions and Availability Zones

The AWS Cloud infrastructure is built around Regions and Availability Zones (AZs). AWS Regions provide multiple, physically separated and isolated Availability Zones which are connected with low latency, high throughput, and highly redundant networking. These Availability Zones provide an easier and more effective way to design and operate applications and databases, making them more highly available, fault tolerant, and scalable than traditional single data center infrastructures or multi-datacenter infrastructures.

Each AWS Region has multiple Availability Zones and data centers. To help ensure the availability and performance of our products, Medidata deploys MCC components across multiple Availability Zones in the same region for fault tolerance and low latency. Availability Zones are connected to each other with fast, private fiber-optic networking, enabling you to easily architect applications that automatically failover between Availability Zones without interruption.
In addition to replicating applications and data across multiple data centers in the same Region using Availability Zones, Medidata also increases redundancy and fault tolerance further by replicating data between geographic Regions, using both private, high speed networking and public internet connections to provide an additional layer of business continuity, or to provide low latency access across the globe.

Medidata utilizes the following AWS regions:

- Northern Virginia. Consists of 6 Availability Zones
- Oregon. Consists of 4 Availability Zones
- Frankfurt, Germany. Consists of 3 Availability Zones

For all MCC components hosted in the US, Availability Zones in the N. Virginia Region serve as the primary Hosting Environment. Medidata has complete control over the application deployment process, including the selection of the Region and Availability Zones being utilized.

Data Backup

Private Cloud

All data in the Houston Data Center operates from a central Storage Management System. This includes both the virtual servers used to operate Medidata services, and all data. This Storage Management System utilizes physical disk redundancy (using RAID 5, and RAID 10) to protect against physical drive failures. The Storage Management System is partitioned into Production and Backup Operations. The Production partition services all server images and data.

Medidata utilizes VMWare's Site Recovery Manager to asynchronously replicate virtual servers and data to Medidata's Frankfurt Data Center (see below). This includes all tiers of virtual machines (web, app and database) with deployed and configured Medidata products, and all data. Recovery at the Frankfurt Data Center is scripted through Site Recovery Manager.

In addition to the replication system, point-in-time backups are maintained, as below:

- 15 minutes – database transaction logs are backed up within the Backup Operations partition of the Storage Management System every 15 minutes. These logs are kept for a minimum of 14 days.
- Every 8 hours snapshots are taken of all production database servers. These snapshots are encrypted prior to storage. The snapshots are stored onsite for 14 days prior to being moved to offsite storage. Total retention time for the snapshots is 90 days (14 days onsite, 76 days offsite)

Amazon Web Services

Medidata Rave Clinical Cloud data is protected using one of 3 different AWS services; Relational Database Service (RDS), or the Amazon Simple Storage Service (S3) or Amazon Elastic File System (EFS).

In production environments, the Relational Database Service (RDS) instance is run as a mirrored pair in different Availability Zones. In the event that the main DB fails, AWS will switch over to the mirror automatically; the application nodes of the product do not typically require a restart. Snapshots of the database server are taken at 5-minute intervals for the length of a backup retention period (generally 35 days), and stored both in another Availability Zone in the Region, and to a secondary Region.
Use of S3 Storage follows a more traditional backup scheme. Database backups are taken at the following schedule, and saved to S3 storage:

- 15 minute transactional
- Daily Differential
- Weekly Full Backup

Applications which require use of a Network File System (NFS) are configured to use Amazon's Elastic File Service. NFS provides availability of stored files across all Availability Zones in the hosting region, and Medidata executes daily backups to a second AWS Region, for disaster recovery.

Disaster Recovery

Medidata recognizes the real possibility of a natural or man-made disaster directly or indirectly impacting the ability of one of its hosting sites to operate. As a means of mitigating that risk, Medidata has developed a Disaster Recovery Plan that describes the method by which Medidata will re-establish, for any Medidata Hosting Environment, business operations utilizing a Recovery Site.

As a second layer, each Primary Hosting Environment is supported by a Disaster Recovery Hosting Environment of similar qualifications

As examples:

- Medidata's Primary Hosting Environment (private cloud) in Houston, Texas, utilizes Medidata's Frankfurt Data Center for disaster recovery. All infrastructure in the Frankfurt Data Center is built following the same Qualification procedures as used in the Houston Data Center.
- Medidata's primary Hosting Environment (public cloud) in Amazon Web Services (AWS) utilizes AWS services to provide High Availability across the Availability Zones (AV's) in the primary AWS Region, and the same AWS services for recoverability in alternate Regions, as appropriate.

Testing and executing the Disaster Recovery Plan helps to verify that the recovery procedures work as intended and that the supporting documentation is accurate and current. Testing also provides an opportunity to identify any improvements needed to the recovery procedures or documentation and to determine whether personnel are adequately prepared to perform their assigned duties.

At minimum, Medidata will schedule and execute recovery tests if change to product or hosting architecture significantly impacts the current disaster recovery solution (data and product backup and recovery systems), with no more than 365 days between tests. If there are changes to Hosting Environment or product architecture, which impact Medidata's Disaster Recovery capabilities, then testing will be conducted to verify those changes.

Testing is conducted in two (2) phases:

Technical Tests will include the technical portion of a DR, performed as a hands-on recovery to the Recovery Site.
Tabletop tests include a review of roles & responsibilities, communication and decision making in declaring and responding to a disaster.
Medidata maintains a Disaster Recovery Testing Summary Report which provide explanation of the most recent testing for each MCC component, including:

- A summary for the testing of components within the MCC Issues identified during the test
- Recommendations for improvement

Disaster Recovery Strategies

Disaster Recovery strategies (and Disaster Recovery Test strategies) are evaluated in 4 tiers:

Highly Available	Products function continuously using components running in *n*+1 regions. If a single region fails, traffic is automatically routed to available regions. No failover necessary.
Failover	Critical components failed over to a secondary region where previously configured and verified.
Recovery	Critical pieces of the platform are relaunched in an available region where data has already been replicated.
Rebuild	Critical pieces of the platform are rebuilt with required functionality, but without data.

Private Cloud Hosting (US)

Normal Operations: Applications are hosted on a virtual server infrastructure (VMWare). Virtual server pools are built across physical servers (Cisco UCS blade servers), spanning server racks, to provide redundancy. Extra physical server resources are allocated to each virtual server pool to allow physical servers to be taken off-line without impacting application performance or availability.

All virtual server images (VMs) and data are stored and run from a central storage management system. The storage system is partitioned to support production and backup operations using separate resources. The storage system uses RAID 5 Flash Drives for databases, and RAID 5 Solid State Drives (SSD) for application and web tiers. (RAID provides redundancy against failure of physical drives in the storage system)

Disaster Recovery Replication: Storage volumes (including both Virtual Servers and Data) are asynchronous, continuously replicated from Medidata's Houston Data Center to the Frankfurt Data Center.

Point-in-time Backups: 1. Database transaction logs are copied to a Network Area Storage (NAS) share using native SQL tools every 15 minutes. The NAS Transaction Log shares are backed up daily and directly to tape and retained for 90 days. 2. The production database snapshots are taken every 8 hours. The blocks of data are stored encrypted, deduplicated, and compressed. Database snapshots are retained for 90 days.

In a Disaster: If a virtual server fails, it can be restarted or replaced with minimal performance or availability impact.

If a physical server experiences performance issues, it can be replaced with minimal performance impact, as each virtual server pool is built to allow extra physical server resources.

If the primary database becomes corrupt or unavailable, the database can be rebuilt from the daily database backup and transaction log which are saved to the backup partition of the storage system.

Targets*: RPO <15 minutes RTO is dependent on size of the database (database restore time)

If the primary hosting environment were to become unavailable, VMWare Site Recovery Manager is used to restore service by promoting replicated copies of servers and data to become primary.

Targets: RPO <4 hours RTO <8 hours. Note these are target recovery times, NOT SLAs. SLAs are defined in customer agreements

Public Cloud Hosting (US)

Normal Operations: The US-East Region (N. Virginia) is used for hosting the primary product instances. Application nodes are deployed to more than 1 Availability Zone (AZ), to provide load balancing and redundancy.

When utilizing S3 storage, backups are taken at the following schedule and saved to S3 storage in a separate Availability Zone:

- 15 minute transactional
- Daily Differential
- Weekly Full Backup

Applications which require use of a Network File System (NFS) are configured to use Amazon's Elastic File Service (EFS). EFS provides availability of stored files across all Availability Zones in the hosting region. Medidata executes daily backups to a second AWS Region for disaster recovery.

In a Disaster: If an application node fails (or single AZ fails), a new application node is deployed to a different AZ (to maintain redundancy). There is no impact on product availability or performance.

If the AZ hosting the S3 or EFS instance becomes unavailable, the the backup AZ will be promoted to primary.

Targets: RPO 0 RTO <15 minutes

If the entire Region become available, then the application would be restored in the alternate Region, and new application nodes would be deployed.

Targets: RPO <24 hours RTO <24 hours

Identity and Access Management

Perfect identity management is tough; very few organizations are able to achieve perfection. In order to solve this problem, Medidata has implemented a "HRIS-As-A-Master" program, which automates provisioning and deprovisioning, reducing errors and omissions to an absolute minimum. Additionally, we review the entire landscape every quarter, with those reviews being audited by third-party, external organizations for maximum consistency.

Access Management begins with the isolation of our corporate and hosting environments, with "no trust" between environments. Access to hosting requires a second set of unique credentials, for which Medidata applies the 'principle of least privilege' to allow engineers to complete the tasks required of their role, but no more.

Identity Management Standards

For our customer facing systems, the password requirements are:

- Eight-character password length
- Contains at least one uppercase letter, one lowercase letter and one number character
- Moderate strength (e.g., no "ILoveCats")
- Must confirm to Electronic Records Electronic Signatures (ERES) requirements for signature
- Rotated every 90 days; 120 day forced rotation
- Cannot be one of the last twenty passwords

Where possible, we encourage the adoption of Multi Factor Authentication (MFA) and are likely to mandate it in the next 12 months for administrators.

Authorization is separated from authentication, because of the nature of how access is provisioned. iMedidata provides authentication into the environment; authorization grants access to the studies themselves. In order to ensure tight coupling to the supported organization, Medidata customers are responsible for managing access to the studies, as they are closer to the sites, and to promote maximum flexibility and response.

Password complexity is eight characters, with three types of complexity – upper vs lower case, symbol, number and alphabetic characters. We rotate passwords every 90 days.

Because we support 50,000+ sites worldwide, this is the commonly accepted baseline which allows varying technologies to access the Medidata platform.

Multi Factor Authentication

Passwords, even if complex and rotated regularly, are not enough, so we offer Multi Factor Authentication (MFA). This allows either an SMS text message, Authy Authenticator or a voice call to provide additional certainty on the identity of the end user.

We are moving closer to a mandatory MFA across all accounts, as soon as the more remote sites are capable of supporting it. We expect as we get closer to the end of 2021, we will shift the majority of end users to use this key security service.

Endpoint Security

Each server has protections over and above at-rest encryption. Each and every host has the Trend Micro suite of tools installed and is fully integrated into Security Incident and Event Monitoring (SIEM) tool, so that systems are dynamically monitored, and any potential issues are responded to promptly; 24 hours a day, 7 days a week, 365 days a year**.**

Malware

Trend Micro has emerged as a leader in the platform-based malware tooling space; providing not only signature based anti-malware software, but also logic that looks for new zero-day type virus, trojans and other malicious code. This is installed on each and every host in the Medidata environment and is monitored centrally by our Security Operations Center.

Log Inspection

The Trend Micro Deep Security Log Inspection module provides the ability to collect and analyze operating systems and application logs for security events, which in turn feeds the Medidata Security Incident and Event Monitoring systems. Log Inspection rules optimize the identification of important security events buried in multiple log entries. These events can be forwarded to a SIEM system or centralized logging server for correlation, reporting, and archiving. The Deep Security Agent will also forward the event information to the Deep Security Manager console, which is managed by our Security Operations Center.

Intrusion Prevention

Intrusion Prevention Systems (IPS) are ideally suited to detect and stop attacks that originate over the network, including those focused on application and operating system vulnerabilities.

Firewall

This host-based firewall approach complements the Palo Alto firewalls at the perimeter of the private cloud. The same host-based firewalls augment the AWS security groups and network access control lists. These firewalls can block or allow certain types of network traffic by creating a barrier between the client and the network. Additionally, the firewall will identify patterns in

network packets that may indicate an attack on clients. Installed on each and every host, they are integrated into the SIEM systems.

System Hardening

Out of the box system configurations are often insecure; excess services, additional processes and extra posts are commonly default in operating systems, web services, database management systems and other components. Using CIS (Center for Internet Security) benchmarks, the most generally recognized secure system baseline available, we harden all our host and network components to Level 1 of that standard, before they are placed into services, by the use of our automated deployment practices.

System Lifecycle Management

Typically, when a component such as an operating system or Java™ is at the end of life, security support, such as patching, stops. Therefore, it's critical to retire those systems before that point. Medidata's policies is to never have an end of life system or component in service; and aggressively manages updates in order to maintain that posture. Independent quarterly assessment ensures oversight on this; and that systems are upgraded well in advance of their retirement dates.

Update Management

The usual industry standard for applying patches supplied by Original Equipment Manufacturers (OEMs) is thirty days. This means that there is a significant window for malicious activity taking advantage of known and published security defects. At Medidata, we target patching as soon as possible, with the internal target of seven days. Often, it is even less than that for urgent security issues. This maximizes a solid and well defended front against a hostile environment.

Network Security

Data protections must be more than simply encrypting; at Medidata, data is monitored as it flows through the environment. Each packet is monitored, inspected and tracked as traffic flows through the cloud environment. All network devices feed the Security Incident and Event Monitoring System, which in turn is monitored by our Security Operations Center for 7x24x365 response.

Firewalls

The latest enterprise class Firewall is installed at every network perimeter and endpoint. For the private cloud, we use Palo Alto, and for AWS we use Trend Micro endpoint firewalls, including a combination of Security Groups, Network Access Control Lists (NACLs) and Web Application Firewalls (WAF). From an inbound perspective, we only allow ports 80 and 443. 80 is solely for the purpose of redirecting to the secure port for ease of our user community.

Intrusion Prevention and Detection

Palo Alto's Intrusion Detection and Prevention Systems (IDS/IPS) is a network security technology built for detecting and preventing vulnerability exploits against a target application or computer; they also have the ability to block threats in addition to detecting them and has become the dominant deployment option for IDS/IPS technologies. We configure all our IDS/IPS solutions to run in "prevention" mode.
The IPS monitors traffic and reports its results to the Security Incident and Event Monitor (SIEM) and the Security Operations Center (SOC) for rapid response to emergent threats.

Malware Protections

Palo Alto also brings unique malware prevention capabilities across the network, endpoint and cloud. The stream-based engine blocks malware in-line, stopping attacks before they can succeed, without impacting performance, providing high-efficacy malware prevention through multiple techniques, including:

- Consistent protection and enforcement across all deployment scenarios.
- Signatures based on payload, not hash, or other easily changed attributes.
- In-line, stream-based detection and prevention of malware hidden within compressed files, web content or other common file types.
- Near real-time updates from the WildFire® threat analysis service, ensuring protection against zero-day malware.

Anti-DDOS

Medidata, like other organizations, saw a significant increase in Distributed Denial of Service (DDOS) attacks in recent months, so in order to proactively defend against this emergent threat, we added additional anti-DDOS components, over and above the protective features in the Palo Alto firewalls and AWS Shield services. Using F5 Silverline, we are able to silence the botnets before they impact the services that we provide; and using AWS's Advanced Shield, we are protected from the most ambitious and aggressive botnet attacks**.**

Mail Filtering

Over 40% of the mail sent in the world is malware, spam and other unwanted mail. Medidata processes over 16 million email messages a month, so even 99.999% effectiveness implies that some of this will get through. In order to protect against that, Medidata uses Proofpoint for mail filtering, spearfishing attack prevention and secure email messaging, reducing this threat to an absolute minimum.

Defect and Vulnerability Management

A security vulnerability is a special class of defect; in that it not only can affect the stability of an environment, but the confidentiality and integrity of the data and the systems that process that data. Our multi-layered program is iterative, comprehensive and aggressive and designed to minimize attack surface, security vulnerabilities and risk to customer data.

Vulnerability Scans

We scan everything, internal and external, across the entire Medidata environment, every month, using Rapid7 Insight Vulnerability Manager. The results are populated into our ticketing system, evaluated for risk, priority and put into the backlog for remediation. Critical vulnerabilities – those with no compensating controls that pose significant risk – are resolved as soon as possible, with an emergency release if necessary. Vulnerabilities rated “high” are fixed within 30 days, “medium” within 180 days and “low” annually.

Internal scans are credentialed in order to thoroughly assess all registries and system files to get in as deep as possible to scour for bugs, misconfigurations and other defects.

Static Source Code Scans

We also scan every application release prior to General Availability, in order to ensure that no software is released with security weaknesses. Static source code scanning is the practice of analyzing software code prior to compilation and deployment.

Static source scans for the entire code base are done monthly, and prior to product release, using BurpSuite Enterprise, WhiteHat or Brakeman. These scans are also credentialed for maximum effectiveness of the tools.

We also put the source code scanning tools directly in the hands of developers. The earlier in a cycle that a defect can be fixed, the less risk and cost is incurred.

Dynamic Scanning

Dynamic code analysis, or scanning of compiled/deployed code, is capable of exposing a subtle flaw or vulnerability too complicated for static analysis alone to reveal and can also be the more expedient method of testing.

At Medidata, we feel that both dynamic and static scanning provides a holistic and more complete testing regime in order to maximize the number of findings in the most scenarios, and therefore reducing overall risk.

Free and Open Source Software

Managing Free and Open Source software licenses is critical; we have over 8000 open source packages, and in order to effectively manage that, we use FOSSA. This tool allows us to determine the proper licensing model in order to prevent risks associated with inappropriate deployment of open source software.

Threat Intelligence

Medidata uses a variety of tools and practices with respect to Threat Intelligence. Our Managed Security Service Provider (MSSP) alerts Medidata to emergent threats; but we also ingest feeds from a number of data sources, so that new attacks are stopped before they impact our systems. We also use Cisco StealthWatch which provides comprehensive visibility into the network traffic and provides advanced threat detection and accelerated threat response using advanced behavioral analytics.

Penetration Testing

Due to the shared nature of the environment, we cannot allow customer penetration testing or vulnerability scanning of our environment**.**

Because our customers expect comfort around how we protect their data, we provide full transparency into our Information Security Penetration Testing program; posting the detailed tests and their results online for perusal by our customer's security experts. To date, no penetration test has succeeded in gaining access to patient data.

Every 90 days (adjusted to correspond to release cycle), we bring a different world class penetration testing entity to attempt to gain access to our environment, testing our monitoring and alert and in general confirming the integrity of the boundaries of data protection. Coalfire, Optiv, BlackHills, DirectDefense, CDW Security, SecurIT360 as well as others all participate in this program.

Red/Blue/Purple Teams

Once a year, we bring two teams, one offensive, and one working with the defensive team to truly wring out the environment. The mandate is simple – get in, however you can. This provides objectivity and assurance that patient data is protected.

Techniques can be as traditional as using fingerprinting and vulnerability assessment tools, or more esoteric techniques such Bluetooth hacking via drones outside a building.

This is a complex, invasive and expensive exercise; but the old maxim of "the more you sweat in peace, the less you bleed in war" means that this investment of time and energy is worthwhile in ensuring that the protections are intact against the most aggressive threats.

Monitoring

Security Incident and Event Monitor

Medidata is a cloud provider, and we like to “practice what we preach”. Where practical, we use cloud services to support the environment, and a key service is SUMOLogic, a system which is tied to every processing and network device in the organization. Consuming over 3 billion security events a month, SUMOLogic processes each activity, login, logout, failed password attempt, API calls and many others to look for attempts to gain or deny access to the environment.

Managed Security Service Provider

Leveraging SUMOLogic, our third-party Managed Security Service Provider “Smartronix “provides 24-hour, 365 day a year monitoring of our systems, and is empowered to disable any external threat, while escalating internally for prompt notification. Using threat intelligence provided by a number of external services, they are proactive in blocking the problems, before they become problems.

Smartronix, a Department of Defense support organization provides around the clock monitoring of all 20,000+ systems in the Medidata universe. With Security Operations Centers around the world, complete global coverage and around the clock response ensures proper support. This service is also supported by additional in-house Network and Security Operations Centers which manage the internal aspects of our NIST 800-53 compliant Incident Response program.

Real-time Configuration Management

Medidata uses Amazon Trusted Advisor in order to provide best practice to the AWS configuration. In order to enhance that, Medidata leverages Palo Alto Prisma Cloud, formerly known as Redlock. Prisma Cloud provides real-time measurement to standards such as NIST 800-53, FedRamp, PCI DSS, HIPAA and other baselines. This allows for assurance that the configuration and management of the hosted systems are properly managed, and as secure as they possibly can be.

Independent Testing and Review

Security is only as good as an objective observer says it is. So, in addition to customer and regulatory audits, Medidata works with multiple independent entities to constantly assess the state of the control environment. We regularly rotate entities to prevent complacency and ensure we bring objectivity, cutting edge techniques and emergent technologies to bear with the goal of ensuring the covered data of our customers and their patients are properly protected.

SOC 1

Medidata published its first SOC 1 Type 1 report for our Rave Site Payments application in 2017, and the first Type 2 in 2018. SOC 1 Type 2 reports are examination engagements performed by a service auditor (CPA) in accordance with Statement on Standards for Attestation Engagements (SSAE) 16, *Reporting on Controls at a Service Organization,* to report on the suitability of the design of the controls at a service organization that are likely to be relevant to an audit of a user entity's financial statements.

SOC 2+

Medidata publishes a Service Organization Controls 2 (SOC 2+) report. This audit is the replacement of the Statement on Auditing Standards No. 70 (SAS 70) Type II report. The audit for this report is conducted in accordance with the Statement on Standards for Attestation Engagements No. 16 (SSAE 16) and the International Standards for Assurance Engagements No. 3402 (ISAE 3402) professional standards. This dual-standard report can meet a broad range of auditing requirements for U.S. and international auditing bodies. The SOC 2 report attests that Medidata data center control objectives are appropriately designed and that the individual controls defined to safeguard customer data are operating effectively. Our commitment to the SOC 2 report is ongoing, and we plan to continue our process of periodic audits.

In addition, Medidata obtains the SOC 2+ with both the information security and privacy trust principles, reinforcing the governance program with a homogenous set of controls around quality, privacy, and information security.

PCI DSS Service Provider

The Payment Card Industry Data Security Standard (PCI/DSS) was created to standardize controls around cardholder data to reduce credit card fraud. Validation of compliance is performed annually by our trained and certified Internal Security Assessor (ISA). We also apply it to any financial processing related information that is used as part of our Payments offering.

ISO/IEC 27001:2013

ISO 27001:2013 is a security management standard that specifies security management best practices and comprehensive security controls following the ISO 27002 best practice guidance. This is a widely recognized international security standard in which Medidata clients showed significant interest.
Certification in the standard requires us to:

- Systematically evaluate our information security risks, taking into account the impact of company threats and vulnerabilities
- Design and implement a comprehensive suite of information security controls and other forms of risk management to address company and architecture security risks
- Adopt an overarching management process to ensure that the information security controls meet our information security needs on an ongoing basis

The key to certification under this standard is the effective management of a rigorous security program. The Information Security Management System (ISMS) required under this standard defines how we continually manage security in a holistic, comprehensive way. The ISO 27001:2013 certification is specifically focused on the Medidata ISMS and measures how our internal processes follow the ISO standard. Certification means a third party accredited independent auditor has performed an assessment of our processes and controls and confirms we are operating in alignment with the comprehensive ISO 27001:2013 certification standard.

ISO/IEC 27018:2014

ISO/IEC 27018:2014 is a security management standard that specifies security management best practices and comprehensive security controls in the context of Privacy Information in a cloud environment. This is a widely recognized international security standard in which Medidata clients also show significant interest.

This standard complements ISO/IEC 27001:2013 and other security frameworks in order to maintain effective management of privacy related information. Like ISO/IEC 27001:2013, ISO/IEC 27018:2014 certification means a third party accredited independent auditor has performed an assessment of our processes and controls and confirms we are operating in alignment with the comprehensive ISO 27018:2014 certification standard.

ISO/IEC 27701:2019

ISO/IEC 27701:2019 is a privacy extension for an existing Information Security Management System; including the adoption of a Privacy Information Management Systems (PIMS). This is a widely recognized international security standard which is as close to a GDPR certification that can be currently obtained. This standard complements ISO/IEC 27001:2013 and other security frameworks in order to maintain effective management of privacy related information. Like ISO/IEC 27001:2013, ISO/IEC 27701:2019 certification means a third party accredited independent auditor has performed an assessment of our processes and controls and confirms we are operating in alignment with the comprehensive ISO 27701:2019 certification standard.

FISMA

Medidata enables U.S. government agency customers to achieve and sustain compliance with the Federal Information Security Management Act (FISMA). FISMA requires federal agencies to develop, document, and implement an information security system for its data and infrastructure based on the NIST (National Institute of Standards and Technology Special Publication) 800-53, Revision 4 standard. FISMA requires Medidata to implement and operate an extensive set of security configurations and controls. This includes documenting the management, operational and technical processes used to secure the physical and virtual infrastructure, as well as the third-party audit of the established processes and controls. Medidata is evaluated every year to maintain our FISMA compliance for Software as a Service and has been awarded an Authority to Operate by a number of US Government agencies.

Privacy Shield

We appreciate that our customers and partners may have questions about the July 16, 2020, European Court of Justice (the ECJ) ruling in a case examining transfers of personal data from the EU to the US under the EU-US Privacy Shield.

What will it mean in practice?

- First and foremost, personal data transfers from the EU to the US using Medidata's services remain secure and compliant with EU requirements. The ECJ's ruling does not change data flows for our services: your use of our US-based commercial cloud services remains in compliance with the ECJ's ruling.
- For years, Medidata has provided our customers with overlapping protections under both the Standard Contractual Clauses (SCCs) and Privacy Shield frameworks for data transfers. Although the ECJ's ruling invalidated the use of Privacy Shield moving forward, transfers based on the SCCs remain valid. Our customers are already protected under the SCCs in place in our standard data protection agreements.
- We note that in addition to the SCCs and the now-invalided Privacy Shield mechanism, informed consent is also a legal basis under the GDPR for the transfer of clinical data to the US for data processing and submission to regulatory authorities. Medidata

recommends that sponsors and CROs consider the use of informed consent documents that disclose clinical data flows from the EU to the US.

- Medidata will closely follow any further guidance from EU data protection authorities and the European Data Protection Board.

FIPS 140-2

The Federal Information Processing Standard (FIPS) Publication 140-2 is a U.S. government security standard that specifies the security requirements for cryptographic modules protecting sensitive information. To support customers with FIPS 140-2 requirements, Medidata Private Cloud VPN endpoints and TLS-terminating load balancers in Medidata (U.S.) operate using FIPS 140-2 validated algorithms. Operating in FIPS-140-2 compliance mode does require comparable capabilities at the user browser side of the connection. While we do not employ FIPS 140-2 certified hardware, we do use the comparable make and model with fully approved FIPS 140-2 software.

HIPAA

For our Rave Commercial Imaging and Quantum Real World Evidence platforms, Medidata enables covered entities and their business associates subject to the U.S. Health Insurance Portability and Accountability Act (HIPAA) to leverage the secure. Medidata environment to process, maintain, and store protected health information.

Conclusion

In closing, we truly believe that Medidata's Information Security program, in conjunction with our Data Privacy and Global Compliance functions, serve as a key differentiator toward ensuring the right patient gets the right drug at the right time.

Please see for yourself by visiting our Trust Center at https://www.medidata.com/trust.

For any questions or clarifications, feel free to reach out to asksecurity@medidata.com.

Appendix 1: Medidata Products by Hosting Environment

Medidata Product	Hosting Environment
Classic Rave	Private Cloud (HDC)
Cloud Administration	Public Cloud (AWS)
iMedidata	Public Cloud (AWS)
Medidata Designer	Public Cloud (AWS)
Medidata Detect (formerly Rave CSA)	Public Cloud (AWS)
Medidata Issue Management	Public Cloud (AWS)
Medidata Patient Profiles	Public Cloud (AWS)
Medidata Remote Source Review	Public Cloud (AWS)
Medidata Risk Management (formerly RACT)	Public Cloud (AWS)
Medidata Site Monitoring	Public Cloud (AWS)
MEDS Extractor	Public Cloud (AWS)
MEDS Perform	Public Cloud (AWS)
MEDS Reporter	Public Cloud (AWS)
myMedidata	Public Cloud (AWS)
Rave Archive	Public Cloud (AWS)
Rave Batch Uploader	Public Cloud (AWS)
Rave Coder	Public Cloud (AWS)

Rave CTMS (Clinical Trial Management System)	Public Cloud (AWS)
Rave Design Optimizer	Public Cloud (AWS)
Rave eCOA	Public Cloud (AWS)
Rave eConsent	Private Cloud (HDC)
Rave EDC (Electronic Data Capture)	Private Cloud (HDC) + Public Cloud (AWS)
Rave eTMF (Electronic Trial Master File)	Public Cloud (HDC)
Rave Grants Manager Contracting and Rave Grants Manager Planning	Public Cloud (AWS)
Rave Imaging	Public Cloud (AWS)
Rave Omics	Public Cloud (AWS)
Rave RCM (Regulated Content Management)	Public Cloud (AWS)
Rave RTSM (Randomization and Trial Supply Management)	Public Cloud (AWS)
Rave Safety Gateway	Private Cloud (HDC)
Rave Site Payments	Public Cloud (AWS)
Rave SOP Management	Public Cloud (AWS)
Rave Trial Assurance	Public Cloud (AWS)
Rave TSDV (Targeted Source Data Verification)	Public Cloud (AWS)
Rave Virtual Trials	Public Cloud (AWS)

Rave Wearable Sensors	Public Cloud (AWS)
Study Management	Public Cloud (AWS)
Sensor Cloud	Public Cloud (AWS)
Site Cloud: End of Study	Public Cloud (AWS)