This content is now outdated.

Please see https://tidepool.atlassian.net/wiki/spaces/PUBSEC/overview for the latest security, privacy, and regulatory documentation.

Tidepool Security, Privacy, and Regulatory Detailed Information

Last updated December 10, 2020

Contact information:

- Howard Look, CEO, CPO and CSO, howard@tidepool.org
- Ben Derr, Security Engineer, ben@tidepool.org
- Christopher Snider, Community and Clinic Success Manager, christopher@tidepool.org

Security - security@tidepool.org
Privacy - privacy@tidepool.org
Legal - legal@tidepool.org

This document has been built to answer and address technical questions related to Tidepool's software related to architecture, infrastructure, security, as well as privacy and regulatory practices. As an inquiring clinic, hospital, or health system, please use this document as your initial reference for your Security Questionnaire or IT audit.

If you have been sent this document in response to a request to fill out a Security Questionnaire or IT audit, please take time to review the information provided. If your audit or questionnaire has a question that is not answered by the information provided, please let us know and we will (1) answer your question and (2) update this document so future inquiring clinics, hospitals, and health systems will benefit from your thoroughness.

If you have any questions about the information provided in this document, please contact security@tidepool.org.

Table of Contents

Overview - What is Tidepool?

Tidepool Software

Tidepool provides software to help import, manage, store and gain insights from diabetes device data. We provide our software and data hosting services for free to end users and clinicians. Tidepool's currently deployed applications include:

- Tidepool Web, a hosted web application for visualizing and gaining insights from diabetes data. Tidepool for web is accessed via https://app.tidepool.org.
- Tidepool Uploader, an installable Mac and PC application that allows end users and clinicians to upload data from diabetes devices such as insulin pumps, continuous glucose monitors and blood glucose meters.
- Tidepool Mobile, mobile applications for iOS and Android that allow end users to add context to their diabetes data. iOS users can also use Tidepool Mobile to automatically upload Dexcom CGM data from HealthKit. Clinics typically do not use Tidepool Mobile.

End users and clinics can sign up for a new account at https://tidepool.org/signup. When logged in as a clinician, the Tidepool Uploader presents a user experience tailored to quickly creating and retrieving patient profiles in a clinic. See "Accounts and Data Sharing Permissions" below for more information about clinic-created patient custodial accounts and end-user created accounts.

More information about Tidepool's software experience can be found at https://tidepool.org.

Tidepool, the Company

Tidepool is a California non-profit organization, and is also a 501(c)(3) non-profit under federal law:

- Federal EIN: 46-2302287
- California SOS Corp. Number: C3503633
- DUNS Number: 079928746
- FDA Owner Operator Number: 10050969

Office (for in-person visits):

Tidepool
Venrock Building
3340 Hillview Ave
Palo Alto, CA 94304

Mailing Address:

Tidepool
555 Bryant St., #429
Palo Alto, CA 94301

Tidepool, the Open Source Project

Tidepool is an open source project. All of our software source code is freely available at github.com/tidepool-org. Our software source code is made available without restriction, using the permissive BSD 2-clause open source license: https://opensource.org/licenses/BSD-2-Clause

---

HIPAA Policies and Business Associates Agreements (BAAs)

HIPAA Training and Compliance

Tidepool complies with all HIPAA security, privacy and breach notification rules. All employees and independent contractors are required to review HIPAA training materials and to undergo a HIPAA security audit for all computers and mobile devices that access Tidepool's servers on an annual basis.

Executing a BAA with Tidepool

At this time, Tidepool is not a covered entity under HIPAA. However, your institution may be a covered entity under HIPAA (which is likely why you are reading this document). Tidepool enters into BAAs with health care systems and other covered entities by request. Please contact legal@tidepool.org to discuss completing a BAA with Tidepool.

Our standard BAA (PDF, DOC) is available for review and use.

Tidepool's BAAs with Subcontractors

Tidepool enters into BAAs with our underlying technology providers. These currently include:

- Amazon, for Amazon Web Services
- Google, for G Suite (email, drive and Google docs) and Google Compute Platform
- Rollbar, for HIPAA-compliant client-side logging, monitoring and analysis
- Sumo Logic, for HIPAA-compliant server-side logging. monitoring and analysis
- ZenDesk, customer support ticketing and knowledge base. Found at support.tidepool.org
- MongoDB Atlas for database software, hosting and services

Legal, Privacy Policy, and Terms of Use

Tidepool's general counsel is Kurt Taylor of Wilson, Marshall and Taylor, Palo Alto, CA. Privacy Policy and Terms of Use development was provided by Pillsbury Winthrop Shaw Pittman. Tidepool legal can be reached at legal@tidepool.org.

When creating a new account with Tidepool (identified by a unique email), users, including clinics, need to agree to Tidepool's Privacy Policy and Terms of Use. Full text of these agreements can be found here:

- Tidepool Privacy Policy found at https://tidepool.org/privacy-policy
- Tidepool Applications Terms of Use found at https://tidepool.org/terms-of-use

---

Regulatory

US FDA

Tidepool is registered and listed with the FDA as follows:

- FDA Entity Registration: Tidepool Project, Registration # 3012128418
- Tidepool Platform and Tidepool Uploader
    - Classification: Medical Data Display System
    - Product Code: OUG
    - Regulation: 880.6310
- Tidepool Web, and Tidepool Mobile (formerly known as Blip and Blip Notes)
    - Classification: Continuous Glucose Monitor Data Management System
    - Product Code PHV
    - Regulation 862.2120

More detailed registration and listing documentation can be found here: DOC-0006 US FDA Regulatory Strategy – Tidepool Software. Tidepool's entire regulatory quality system is openly available here.

Tidepool is one of nine participant companies in the FDA Pre-Certification Pilot Program.

International

Tidepool is not marketed outside of the US. Tidepool does not currently have CE marking or ISO 13485 certification. Use of Tidepool's software outside of the US may not comply with local law or with our Terms of Use or Privacy Policy.

Tidepool's Terms of Use state:

> You understand that to register as a User of Tidepool Apps you must be in the United States. We make no claims that Tidepool Apps are accessible or appropriate outside of the United States. Access to Tidepool Apps may not be legal by certain persons or in certain countries. If you access Tidepool Apps from outside the United States, you do so on your own initiative and are responsible for compliance with local laws.

Tidepool's Terms of Use and Privacy Policy are compliant with the requirements of the EU General Data Protection Regulation (GDPR).

System Architecture

Overview

Tidepool's cloud platform is hosted on Amazon Web Services (AWS) with data hosted in MongoDB Atlas. Tidepool has multiple, isolated network and compute environments, used for different purposes, architected and informed by AWS recommendations for HIPAA-compliant services:

- PRD - our production server environment; this is where all end-user data is stored, including PHI covered under HIPAA.
- INT - our integration test environment for 3rd party developers to use as a test sandbox for their applications that access our APIs; this environment may also contain PHI covered under HIPAA.
- Staging - staging environment for final testing prior to deployment to production and integration test environment; this environment is not used for PHI covered under HIPAA.
- QA1 and QA2 - development environment, for day-to-day iteration by our development staff; this environment is not used for PHI covered under HIPAA.

The PRD and INT environments are hosted on HIPAA-compliant AWS instances. (QA/Staging systems are not used with PHI data covered under HIPAA, but are treated as if they do).

The PRD environment features multiple auto-scaling groups and is hosted in multiple availability zones. The production database is hosted in triplicate (primary-secondary-secondary) in multiple availability zones, and is additionally backed up hourly. All data, at all times, is transmitted using secure, industry-standard, encrypted protocols (HTTPS, TLS). All data, at all times, is encrypted at rest (using AES-256) and in transit (using TLS 1.2 or higher). Access to servers is limited and logged. No access to data is permitted to our data hosting provider, MongoDB Atlas.

The overall system architecture is shown below. A more detailed diagram and description of the production environment (PRD) follows.

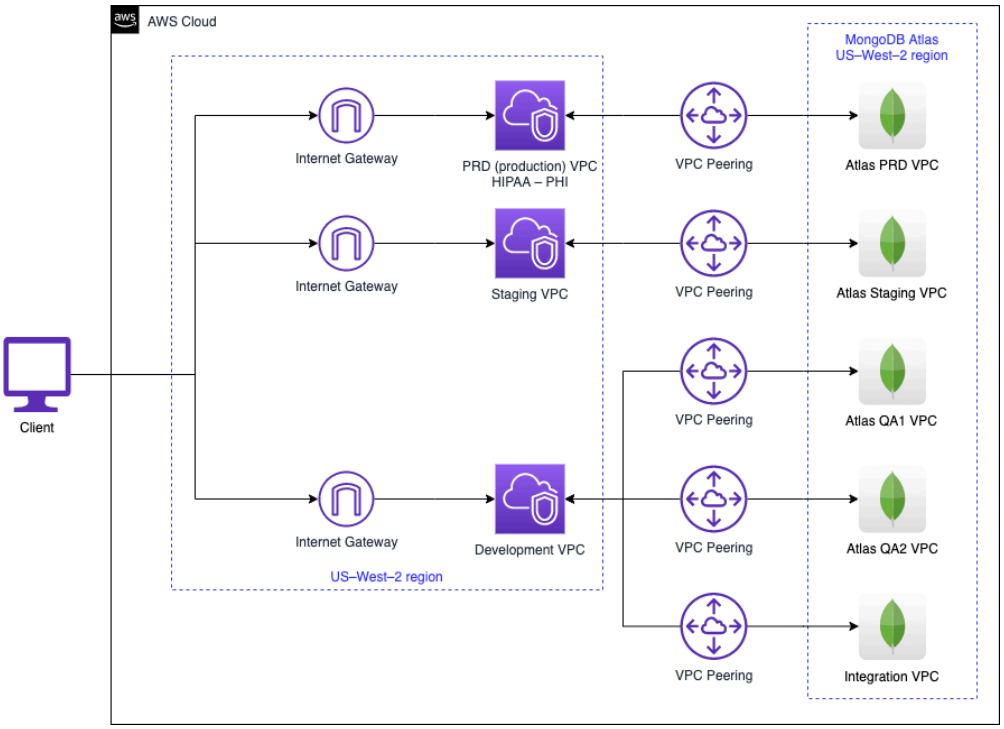

Figure 1: Tidepool Overall AWS System Architecture

The PRD environment is logically isolated in its own Virtual Private Cloud (VPC) from other environments (QA, Staging, Integration). Different environments can access each other only as explicitly permitted via policy and permissions managed in infrastructure as code (e.g. GitOPS)

All of our servers are currently hosted in the US-West 2 Region of AWS. Redundant servers are maintained in multiple availability zones within US-West 2; detailed below.

Server Configuration: Production (PRD) and Integration (INT) Environments

The diagram below shows the Tidepool Production (PRD) environment in detail.

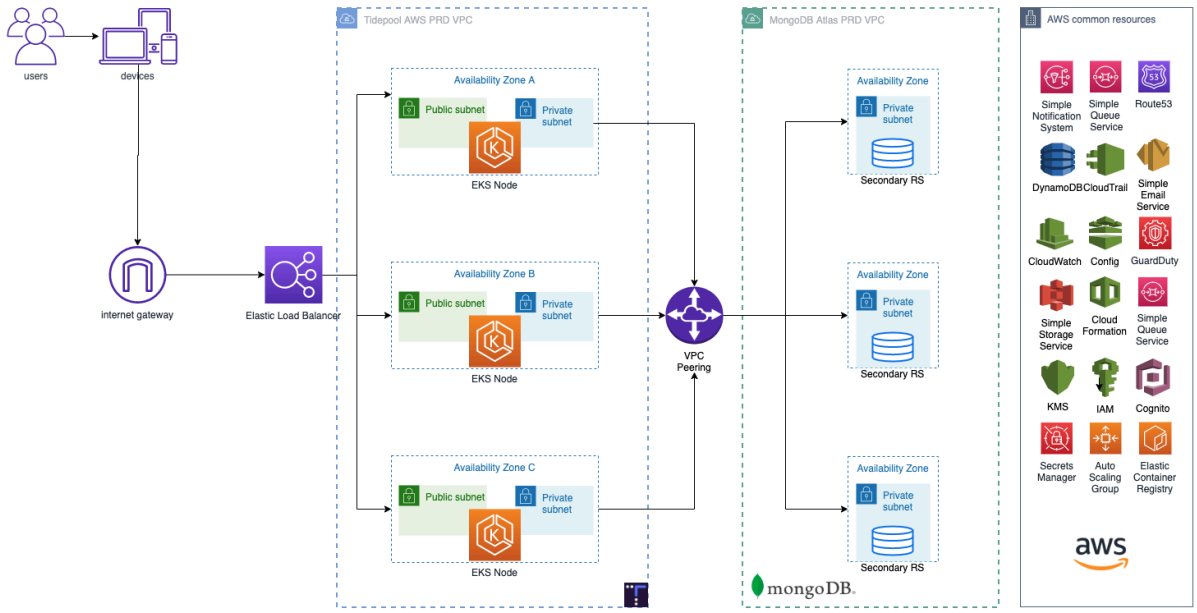

Figure 2: Tidepool Production Environment (PRD)

All PRD and INT servers are hosted on dedicated, HIPAA-compliant AWS EKS Kubernetes clusters.

Our QA and Staging environments are only used internally for Tidepool development, do not host end-user PHI, and are not covered under our HIPAA BAAs, but run in identical and separate environments for consistency in development, testing and data isolation

All user data is hosted in MongoDB instances running inside of MongoDB Atlas. These DBaaS (Database as a Service) instances also run on AWS, but with the infrastructure for the databases handled as a managed service, under the care of MongoDB. MongoDB Atlas handles backups, software upgrades/patching of databases, network security.

All environments also store data in secure private encrypted AWS S3 buckets. There is no access granted to these S3 buckets other than via explicit policy, and only systems hosting the data and Tidepool AWS administrators may access the data.

The PRD environment consists of:

- cluster-production: the Tidepool application, hosted in a kubernetes cluster.
- atlas-prd-db1, atlas-prd-db2 and atlas-prd-db3: the MongoDB Atlas production cluster, configured as a replica set of three servers in a primary-secondary-secondary configuration (MongoDB info), in separate availability zones (more below under MongoDB database)

AWS Services

PRD also relies on these other AWS-provided services, which may be hosted in locations other than US-West 2:

- AWS S3 - AWS Simple Storage Service
    - daily, encrypted backups of MongoDB replica set
    - raw device upload data
    - logs of all access to AWS resources via CloudTrail, CloudWatch, Config
    - Note: S3 is not region-specific. Even if US-West 2 goes down entirely (a catastrophic event that would cause major Internet service disruption), S3 is still available. This would allow Tidepool to spin up a new production environment in a new region in a matter of hours.
- AWS SES - AWS Simple Email Service
    - Used to send all Tidepool service emails (e.g., new account verification, forgot password)
    - Two dedicated SES instances assist in security/logging of email/email monitoring
- AWS DynamoDB
    - Used for saving application metrics (also see below, Logging and Metrics)
- AWS Route53: provides Tidepool DNS
- AWS CloudTrail: records AWS API calls
- AWS CloudWatch: provides monitoring for AWS cloud resources and applications
- AWS CloudFormation: management and provisioning of AWS resources
- AWS Config: resource inventory, configuration history, and configuration change notifications for security and governance
- AWS EKS: managed Kubernetes (k8s) services
- AWS Secrets: handling/storing/providing sensitive configuration data
- AWS GuardDuty: network security monitoring
- Keycloak: Future IMS (Identity Management System), not implemented yet
- AWS ELB: Amazon Elastic Load Balancing
- AWS SNS: Amazon Simple Notification Service - triggered relay of logs and notifications
- AWS ASG: Auto Scaling Groups assist in maintaining and dynamically adjusting desired number of systems/resources

All environments are configured identically to PRD, with these differences:

- Separate networking environment, database and compute
- Reduced resource demands (smaller environment)

All of the instances in each environment are self-contained in AWS VPC's (Virtual Private Cloud) to allow network isolation so a change in Development has no way to affect anything in Production).

There are also a variety of support infrastructure for each environment (load balancers, auto-scaling, public/private networks, monitoring and logging, metrics, etc.). These resources are accessed by the environments as common services in a separate cluster.

Environment Isolation via VPC, Access Control

VPC Peering is used to support the private networking of the AWS and MongoDB environments while maintaining boundaries (though MongoDB is itself running on AWS). AWS VPC peering is used to join the private networks together.

Security Groups and Network Access Control Lists (ACLs), define which ports are exposed and which machines can talk to which other machines via port/machine/destination/subnet masks.

Availability Zones, Redundancy within each VPC, Auto-Scaling Groups

Within each VPC, multiple redundant servers are kept within different Availability Zones. Per AWS, "Each Availability Zone is engineered to be isolated from failures in other Availability Zones, and provide inexpensive, low-latency network connectivity to other zones in the same region." [ref]

Each kubernetes cluster, for example **cluster-production**, is logically grouped within an AWS auto-scaling group (ASG), allowing for easy programmatic expansion when necessary, but are also housed in separate Availability Zones for improved redundancy. Further, nodes within a cluster are in **3** separate Availability Zones for additional resource isolation and redundancy.

Use of Kubernetes

Each kubernetes cluster is managed by a single GitHub configuration repo. A repo can be publicly readable or private. A private repo can also differentiate people by their GitHub identity and assign them read-only access, read-write access, or administrative privileges. Administrative privileges are only granted to Tidepool Ops team members for the repos below.

Within the config repo is a directory that contains the configuration of the services that are shared across the cluster, including configuration of logging, security, and high availability.

Each Kubernetes cluster hosts one or more services. The services include Tidepool environments, and other services that we need for monitoring or sharing our artifacts.

Each Tidepool environment defines an instance of the Tidepool backend services. Each environment is completely independent of the others. Consequently, each environment can run whatever version of the Tidepool services in accordance with the needs of the clients of the environment.

Each Kubernetes environment is backed by storage. That storage consists of Mongo database storage and file storage. That storage can be persistent or ephemeral, in accordance with the needs of the clients of the environment.

Each Kubernetes environment is accessed via one or more DNS aliases.

Use of MongoDB

End-user data is stored primarily in MongoDB, configured in a Primary/Secondary/Secondary "replica set" configuration (more info on MongoDB replica sets). In this configuration, if the primary goes down, the secondaries "vote" on who becomes the new primary. When the old primary comes back, it is told that a "vote" occurred, and it becomes a secondary.

All writes and reads are to the Primary. MongoDB handles making the two secondaries "eventually" consistent. In practice, secondaries are consistent within a second under normal circumstances.

Power, network, hardware and drive failures are to be expected. Services are designed to be fault-tolerant and self-healing where possible.

All filesystems that contain user data, including the MongoDB filesystems and AWS S3 buckets, are encrypted using AWS encryption keys exclusive to Tidepool or MongoDB Atlas and managed internal to AWS.

Network Configuration, Bastion Subnet, Access Control, and Logging

All incoming requests from the Internet arrive via AWS Elastic Load Balancer (ELB) services. The ELB service is unique per environment. Amazon ensures that the AWS ELB service is redundant and fault-tolerant.

Tidepool also uses the AWS NAT Gateway Service to allow internal instances to reach out to the Internet, for example to reach out for time updates (via NTP), but does not allow outside traffic in. The NAT Gateway also provides automatic fault tolerance.

Only strictly necessary ports are configured:

- Ports 80 (http) and 443 (https) on app instances
- Port 123 Network Time Protocol (NTP) on all instances

HTTP port 80 is only used by our marketing web site, tidepool.org. Our app redirects HTTP port 80 to HTTPS port 443. Our end-user and clinic application site, app.tidepool.org, and our api endpoint (api.tidepool.org) only uses HTTPS port 443. HTTP port 80 automatically redirects (via 301 Permanent Redirect) to the same page on HTTPS over port 443. Internal servers with no public-facing ports use self-signed certificates so that all data communication is still encrypted in transit at all times.

Access to individual kubernetes and database nodes for administration is provided via the Kubernetes proxy, which requires AWS IAM credentials using multi factor authentication or via a Bastion host as detailed above.

No other direct remote access is possible to application clusters and all changes are handled via a GitOPS workflow, where everything including the infrastructure is being managed via code.

Some Tidepool EC2 instances are used for development and testing, so Bastion Hosts are retained for this reason, though there is no access to kubernetes clusters directly except through it's API, granted by AWS authentication with Multi Factor Authentication enforced.

Drive Configuration, Encryption at rest, and Backups

All production instances use encrypted, ephemeral storage. All filesystems holding user data are encrypted with encryption keys managed by AWS and controlled by Tidepool.

Each MongoDB database is backed up hourly to MongoDB Atlas. In the extremely unlikely event that all MongoDB instances (in a single environment, in separate availability zones) went offline or were corrupted, the entire production database could be recreated from the backup.

Additionally, backups will be stored externally to AWS S3 storage in an abundance of caution, though MongoDB Atlas is using similar mechanisms of retention.

Software Architecture and Technology Stack

Software Technology Stack

Components of the Tidepool Software Architecture stack are summarized below:

Function	Technology or Service Used
Data storage	MongoDB Atlas, Amazon S3
Server architecture	NodeJS, Go
Client architecture	ReactJS, Flux
Data visualization	D3
Rollbar	Client-side HIPAA-compliant logging
SumoLogic	Server-side HIPAA-compliant logging
Keycloak	Identity Management and Authentication (in development)

All of Tidepool's software source code is freely available for inspection, copying and reuse at github.com/tidepool-org.

Service Endpoints

All Tidepool services are accessed via RESTful URLs at api.tidepool.org. Use of our APIs is documented at developer.tidepool.io and via source code and documentation found at github.com/Tidepool_org. All API calls are RESTful and require a unique session token that is obtained during authentication. Session tokens are 1024 bits and expire after 30 days or upon logout. All communication with APIs is encrypted via TLS/HTTPS. API result bodies are returned as JSON.

Account Creation, Authentication, and Password Protection

Clinics and end users create new accounts using tidepool.org/signup. Username is an email address. Accounts are verified using an email that includes a unique URL based on a unique, 192-bit verification key.

Account passwords must be between 8 and 72 characters and not contain whitespace. No other restrictions are imposed. Passwords are not stored; a unique hash is created using one-way SHA-1 with a private salt.

Keycloak integration is currently under development as a new identity management and authentication provider which will provide advanced security and management functionality, including stronger password policies, secure password reset functionality and potential integration with Enterprise user stores via SAML/SSO.

Authentication is performed using username and password over an encrypted HTTPS connection. A 1024-bit session token is returned to the client. This token must be used for all subsequent API requests.

Accounts and Data Sharing

Clinics may create custodial accounts on behalf of their patients. Upon account creation, the clinic may optionally include an email address for the patient, which causes an email with a secure account verification link to be generated and sent to the patient (end-user).

End-users may also create (non-custodial) accounts outside of clinics. These end-users may select "Share" within the Tidepool web interface and invite another user to view their data via their email address.

Sharing permissions:

- *View:*
    - This user may view my data. For example, end users may invite their clinician or friends to view their data.

- *Upload:*
    - This user has permission to upload data on my behalf. For example, an end user may give a spouse or a clinician upload permissions.

When clinics create custodial accounts for their patients, the clinic is automatically given view and upload permission. If end users create an account outside of the clinic, they may share their data with their clinician, as described above, providing view permission. They also have the choice of allowing upload permission.

Logging and Metrics

User application usage patterns are logged via KISSMetrics. No PHI is stored in these metrics. End-users are identified in the metrics using a one-way cryptographic hash that is distinct and different from their User ID.

Metrics logging is mirrored via AWS DynamoDB. AWS API logging is also logged via CloudTrail. All application logs are stored on encrypted filesystems in VPCs as described above.

Sumo Logic is used to capture, monitor and analyze server-side logs. Rollbar is used to capture, monitor and analyze client-side logs. Both Sumo Logic and Rollbar are under HIPAA BAA with Tidepool for handling of PHI.

Sumo Logic service logs and analysis are available to eight employees who maintain Tidepool's infrastructure. Rollbar logs are available to all developers.

Security Processes, Monitors, Alerts and Access

Uptime, Monitoring and Alerts

Tidepool uses multiple mechanisms for monitoring all environments, including: PagerDuty, Slack and DataDog. An "on-call" rotation schedule is maintained to ensure that there is always a primary and multiple backup employees to respond to potential issues, 24x7.

Additionally, since we are a fully distributed and remote company, Tidepool employs engineers in multiple Time Zones, so an engineer is always available.

Based on Pingdom (legacy uptime monitoring platform) and DataDog reports, Tidepool has maintained 100% user-facing uptime of our production environment over the last year, and over 99.9% uptime since inception (and yes, we just knocked on wood). Individual instances are only taken down momentarily for software installation. User app and API requests continue to be fulfilled by redundant instances.

Persons with Root Server Access, Access to PHI

Only seven Tidepool employees have full administrator access, including production database:

Tapani Otala, Lennart Goedhart, Ben Derr, Derrick Burns, James Raby, Todd Kazakov, and Adin Hodovic.

Other software engineers having software deployment access, but not production database access: Chris McGee and Clint Beacock.

Some end-users may choose to share their Tidepool account data with support@tidepool.org to assist with customer support questions. All full time Tidepool employees (including full time contractors) provide customer support and have access to end user accounts that have been shared in this way.

Security Patch Process

Tidepool servers are evaluated continuously for O/S and supporting software updates. Tidepool's software is typically updated multiple times per week (not all components are

updated, obviously). Urgent patches, though extremely rare, can be tested and deployed within minutes.

Tidepool Employee Password Management, Laptop, and Mobile Security

Tidepool employees are required to use unique, strong passwords and to maintain their passwords in a password manager such as 1Password or LastPass. All Tidepool employees are required to use 2-factor authentication for all Tidepool accounts.

All Tidepool employees are required to use filesystem encryption for their computers, and to have strong passwords and encryption enabled for their mobile devices that access Tidepool accounts.

All SSH key pairs are required to use strong passphrases.

All Tidepool employees undergo annual HIPAA audits of all computers, accounts and mobile devices used for work purposes, including the above requirements. All new employees undergo the same audit.

Infrastructure Security processes and practices

Are the front-end, application, and database components on separate networks?

Yes. Front end applications run in the user's web browser. Back end and database components are on separate, isolated networks (see diagram and description above). There are three "networks"; the "public" network that houses the application instances and NAT Gateway; the "private" network that houses the database; and the "bastion" network in the OPS environment that houses the single bastion instance.

Where and how does data enter and leave the system?

Via RESTful URLs and HTTPS/TLS over the internet.

Does Tidepool host its own servers?

No. Tidepool operates the system, but all servers are hosted via Amazon Web Services (AWS). PRD (production) and INT (integration) instances are on dedicated, HIPAA-compliant AWS instances. User data is hosted in AWS S3 storage and MongoDB Atlas, a managed database service.

Are network environments and virtual instances designed and configured to restrict and monitor traffic between trusted and untrusted connections? If yes, how often is the configuration reviewed?

Yes, see Figure 2 above. Configuration is continually reviewed by our development and operations staff.

Is data from health systems logically or physically separated within the Tidepool environment?

Data from individual clinics or health systems is not kept physically separated, but there is no risk of commingling different clinics since each user is protected by a distinct user ID. Data is logically separated by user ID.

Data Protection and Privacy processes and policies

Is there encryption in transit for all communications and connections including client to server and server to server?

Yes.

Where are Tidepool's servers located?

Tidepool uses dedicated, HIPAA-compliant AWS instances hosted in the `us-west-2` region (Oregon). Tidepool also uses DBaaS (Database as a Service) from MongoDB Atlas, which also resides in the `us-west-2` region.

Does Tidepool support the use of digital certificates?

Yes, *.tidepool.org uses digital certificates signed by AWS and LetsEncrypt for externally-facing connections. Internal networks may use self-signed certificates. All traffic, internal and external, is encrypted at all times.

Does Tidepool require remote health system connectivity for support?

Tidepool does not require remote connectivity from the health system, but the health system will need to be able to access Tidepool's servers over the internet via the web.

Is there an active and documented process for ensuring that inactive accounts are reviewed and disabled if no longer required?

Yes, Tidepool has an active and documented process for removing inactive Tidepool employee accounts. Accounts may be deleted upon request to support@tidepool.org.

Is the local administrator renamed to a non-obvious account name?

We do not allow SSH access from root; only accounts explicitly named in our configuration can access our hosts. Login (via password or SSH) is not allowed for the root account and as a rule SSH is disabled on all production systems and enforced by configuration

Are all unnecessary ports and services disabled on each server supporting Tidepool?

Yes, the network is protected using Security Groups and Network ACLs. Tidepool uses a standard micro distribution of Linux called Alpine which is secured by default.

Are the operating system & application patches tested and approved for implementation?

Yes, Tidepool has a robust, FDA-compliant quality system and follows the software development and testing process described here.

How long does Tidepool retain application and user data and logging?

Per Tidepool's Privacy Policy, Tidepool will retain your account and related information on your behalf as long as needed to support your use of the Tidepool Apps, for necessary backup purposes and comply as necessary with our legal obligations, resolve disputes, establish legal defenses, conduct audits, pursue legitimate business purposes, enforce our agreements and comply with applicable laws. We may delete your account due to inactivity, but we will notify you by email prior to doing so and give you a reasonable opportunity to either transfer your information or begin active use of your account again.

Practically speaking, this means Tidepool does not delete user data unless requested to. At this time, Tidepool retains operation logging and audit trails indefinitely.

Is the application compatible with antivirus software AND will antivirus be implemented?

Tidepool's client application runs in a web browser, and is therefore compatible with any locally-installed antivirus software. Some antivirus programs may complain upon installation of

the Tidepool Uploader device drivers; Tidepool is registered with Symantec, McAfee and Kaspersky.

Tidepool's servers currently run a micro distribution on Linux called Alpine. antivirus software is not installed on these Linux servers. Since all Tidepool services run behind the firewall and only application-specific ports and APIs are open, the risk of a virus is extremely unlikely.

How will Tidepool report/communicate application security vulnerabilities to our health system?

Tidepool will report all security breaches consistent with HIPAA guidelines. Vulnerabilities will be reported to your administrative contact if they have an impact on your health system.

Is there encryption at rest employed for servers and/or databases?

Yes, see above.

Does Tidepool provide encryption at rest for end user devices including laptops, tablets, and mobile phones?

Yes, all Tidepool employee laptops devices are encrypted, and regular audits are conducted to ensure ongoing compliance.

Does Tidepool provide encryption in transit for all communications and connections including client to server and server to server?

Yes.

Are data input and output integrity routines (i.e., reconciliation and edit checks) implemented for application interfaces and databases to prevent manual or systematic processing errors or corruption of data?

Yes, platform data APIs perform data validity checks.

Does Tidepool's web application use any direct reference objects?

No, data is transmitted to the web application as JSON streams over RESTful APIs.

In Tidepool's web application, are all unique identifiers indexed and referenced in the code, instead of displaying the unique id's?

This is not necessary. API references include visible references to UserID, but all API calls require a valid session authentication token which is not exposed in the URL.

Are session IDs exposed in the URL?
No, session IDs are only included in the encrypted HTTPS header and are not exposed in the URL

Do session IDs timeout?

Are session IDs invalidated during logout?

Are session IDs rotated after successful login?
New session tokens are generated upon login and invalidated on logout. A session token remains valid for 30 days if the user remains logged in.

Does Tidepool use or store a Credit Card Number, Social Security Number, or other personal identifier as a data element or alternate key to the application's database?
No.

Does Tidepool perform ongoing security vulnerability and penetration testing?
Yes. Tidepool relies on AWS' ongoing security vulnerability and penetration testing. In addition, Tidepool runs the Tidepool Responsible Disclosure Program. Under this program, outside security researchers test Tidepool on a continuous, ongoing basis. To date, dozens of security researchers have reported issues, all minor, all of which have been fixed or mitigated.

Does Tidepool have any antivirus exceptions?
No. Tidepool software is whitelisted with Symantec, McAfee, and Kaspersky.

Are there any issues running Tidepool software on encrypted hard drives?
No issues. Tidepool employees do this all the time.

In addition, all of Tidepool's production software is open source, and is therefore available for open inspection and comment. See: https://github.com/tidepool-org .

Is Tidepool user data de-identified?
User data is mapped via a unique identifier to data, so it would not qualify as de-identified data.

From a context perspective, a Tidepool user has access to their own data only, unless access is granted by another user. Clinicians have access to the data they upload and any data the patient uploads only, unless data has been explicitly shared with them by another user.

The Tidepool application is able to map access to specific users and relationships with clinicians in the system, but no user has the ability at the application level to access another user's data without being granted that right. This relationship mapping is maintained internally via a system-generated hash using a salt. This unique mapping allows user records to be associated with clinicians.

Any data donated by users to the Tidepool Big Data Donation Project is completely de-identified. Donating data to Tidepool Big Data Donation Project is completely optional and users may opt-in to donate their data.

Operational Processes and policies

Does Tidepool have a formal IT Security team and incident response program?

Yes, Tidepool has a dedicated Security Engineer, a back-end/ops engineering team and documented on-call rotation for incident response. Additionally, Tidepool employs engineers in multiple time zones since we are geographically distributed, so an engineer is nearly always awake and available.

Is there a documented and tested backup process?

Yes. Production data is replicated to secondary instances of MongoDB in a primary/secondary/secondary configuration, hourly backups in MongoDB Atlas and nightly push to S3. Applications can be rebuilt and re-deployed based on our public source code and processes found at https://github.com/tidepool-org. The backup and restore process has been validated and is tested quarterly.

Are projections of future capacity requirements made to mitigate the risk of system overload?

Yes. Tidepool has recently moved (2019) to using more cloud native technologies like kubernetes and the DBaaS MongoDB Atlas which will allow us to scale rapidly for the foreseeable future in response to growth and change.

Compliance and Risk Management processes and policies

Is regulatory compliance (HIPAA) supported by Tidepool?
Yes, Tidepool meets all HIPAA and HITECH requirements.

Does Tidepool perform audit logging of all user activities and ensure logs cannot be modified or tampered with?
Access logs to instances containing PHI are maintained via operating system logging mechanisms. Monitoring, audit controls and system activity review is documented and complies with 45 CFR 164.308(a)(5)(ii)(C), 45 CFR 164.312(b) and 45 CFR 164.308(a)(1)(ii)(D).

Tidepool stores encrypted logs on AWS and in Sumo Logic. Sumo Logic provides monitoring, reporting and alerting

Can Tidepool provide retention of logs that meet HIPAA and other legal and regulatory requirements?
Yes, Tidepool implements administrative safeguards compliant with 45 CFR 164.308(a)(1) and has addressable safeguards compliant with 45 CFR 164.308(a)(3).

Are the availability, quality, and adequate capacity and resources planned, prepared, and measured to deliver the required system performance in accordance with legal, statutory, and regulatory compliance obligations?
Yes.

Identity and Access Management processes and policies

What method does Tidepool use to authenticate end users? to the system and/or application?
User accounts are secured via email address and password.

Email addresses are validated upon account creation by sending an email with a secure (per user hash) URL which must be clicked on prior to the account being verified.

Does Tidepool integrate with OAuth, Active Directory, LDAP or other site-wide authentication?

Tidepool currently provides authentication via email and password. A major upgrade to our Identity Management System is currently under development/integration using **Keycloak**, which will provide additional options like SSO/SAML, OAuth, social login, 2FA/MFA (Two factor or Multi-factor authentication).

What are Tidepool's end user password management policies?

Users of Tidepool are required to have a password of between 8 and 72 characters with no whitespace. **Keycloak** (in development) will provide additional password policy capabilities including 2-factor authentication.

If an end user forgets their password, they can select "Forgot my password," which sends an email to the account of record that includes a reset link with a unique key.

Does Tidepool implement an inactivity timeout feature?

Tidepool does not implement timeout activities. These can be implemented via local workstation policy.

Does Tidepool provide role-based access control for all users?

Yes, role-based account permissions and access control include:

- Edit: Ability to edit data or notes.
- Upload: Ability to upload diabetes data but not user data.
- View: Ability to view data (must be granted by the owner of the data).
- Notes: Ability to attach a note to a user's data.
- Owner: Total control of the account, including metadata.
- Custodian: Permission to access a custodian account (but not change its password)

When an end user creates their own account and selects "Share", they can invite another user to view their data and optionally allow them to upload data on their behalf.

How is administrative access to servers and databases managed?

Access to the administrative console for Amazon Web Services (AWS) is limited to eight employees. Login access to the AWS console uses two-factor authentication. SSH key pairs use strong passphrases and no host can be accessed directly except for a Bastion host.

Application and Database Security policies and practices

Can Tidepool be accessed from a mobile device?

Tidepool's web experience, and the Tidepool Uploader can only be fully accessed from the Chrome web browser on a Mac or PC. Tidepool Web can be accessed with other browsers, including mobile devices, but functionality is limited to non-visualization screens.

Tidepool Mobile (formerly known as Blip Notes) is available for iOS and Android devices. It allows end users to enter contextual information about their diabetes, such as notes about exercise or illness. Tidepool Mobile for iOS can also be used to upload diabetes device data using Apple Health. Tidepool Mobile is typically not used by Clinicians.

Does Tidepool allow users to download, export or print their own data?

Yes.

Does Tidepool have input validation controls to prevent malicious or invalid data from being entered?

All APIs for putting data into the system are secured via authenticated login token and HTTPS/TLS. Email address and password are required to login and receive a token.

Diabetes device data that is uploaded through the Tidepool Uploader is parsed and validated based on device drivers written in compliance with the device maker specifications and/or Open standards.

Have the global database security options and environmental variables been adequately configured to provide a secure operating environment?

These may include such items as:

- **Table space configuration**
- **Remote execution of stored procedures**
- **Query timeout settings**
- **Buffer cache size**
- **Recovery options**
- **Encryption**
- **Initialization parameters**

Yes. Our MongoDB and application servers are in private VPC’s with limited access. All data is encrypted in transit and at rest. All filesystems with data and application software are encrypted.

Will data be stored outside of the U.S.?

No, Tidepool does not currently store production data outside of the U.S.

Tidepool has employees and subcontractors that work remotely from around the world. Employees authenticate using two-factor authentication from wherever they are.

Currently, data for Tidepool applications is only currently stored on AWS servers hosted in the Pacific Northwest region (Oregon, `us-west-2`).

How does Tidepool manage server access logs?

Access logs to instances containing PHI are maintained via operating system logging mechanisms. Monitoring, audit controls and system activity review is documented and complies with 45 CFR 164.308(a)(5)(ii)(C), 45 CFR 164.312(b) and 45 CFR 164.308(a)(1)(ii)(D).

Does Tidepool use virtual machines?

Yes. Tidepool uses Amazon Machine Images (AMIs) built by us programmatically for EC2 as well as a containers running inside of Kubernetes on AWS managed virtual machines.

Does Tidepool use a reliable and mutually agreed upon external time source used to synchronize the system clocks of all relevant information processing systems to facilitate tracing and reconstitution of activity timelines?

Yes, all Tidepool servers sync to nist.gov using Network Time Protocol (NTP) or to Amazon Web Services NTP sources that do the same.

Are production and non-production environments separated to prevent unauthorized access or changes to information assets?

Yes, see the main document for detail on our Production, Integration, Staging and Development environments.

About Custodial Accounts and the Tidepool Data Donation Project

Does Tidepool allow end users to donate their data for research purposes?

Yes. Tidepool operates the Tidepool Big Data Donation Project. This allows end users to opt in to donating their anonymized data for research and innovation uses.

How does the Tidepool Big Data Donation Project work with clinical accounts?

If a clinic enrolls a patient in-clinic, they create a Custodial Account. During the new patient creation process, the clinic can OPTIONALLY enter an email address for the patient, which will send the patient an email and allow them to sign up for Tidepool and claim their account. If they do, the clinic still has access to see data in the account, as if the patient had created the account themselves at home and shared their data with the clinic.

If the clinic does NOT enter a patient email email, the account will remain under the control of the clinic, and the patient will not have access to the data from home, nor be able to upload their own data from home. (In general, Tidepool discourages this - patients really like having access to their data from home. In addition, some clinics have >70% of their patients uploading from home, which saves a ton of time/effort in the clinic!)

For Custodial Accounts that the clinic manages, there is no option to donate that data and no data from those accounts is donated to the Tidepool Big Data Donation Project, which means no anonymized data will make it to third parties.