Monitoring Accounts for School Safety:
AI, Data, & Technology

bark FOR SCHOOLS

Introduction
to Bark for Schools

Students growing up in the 21st century have witnessed the rise of incredible digital technologies. This progress has changed the way students navigate the world, both in positive and negative ways.

On the one hand, they're able to collaborate on projects and access more information than ever before. But they're also better equipped to abuse digital technologies — as well as fall victim to abuse. In 2018 alone, Bark for Schools detected more than 3.5 million concerning situations on school-issued accounts, including:

- **853,185** discussions of using or selling illicit drugs
- **100,284** expressions of self-harm, cutting, or suicidal ideation
- **51** potential acts of violence in schools

Digital technologies have also accelerated the rate at which language evolves, as their development compels corresponding developments in speech and writing. Artificial intelligence and machine-learning techniques provide a way for us to keep pace with these developments, which is especially important when it comes to protecting students from the worst of what exists online. Bark — and the Bark for Schools initiative — is pioneering natural language processing systems to help keep kids and students safe in the digital age.

Bark for Schools monitors millions of online activities every day. This is a massive amount of data to process, and in an effort to provide transparency about our data collection and analysis practices, we have compiled that information into a comprehensive resource.

In addition to account monitoring, we offer a web filtering feature that is not detailed here. This white paper will explain how Bark for Schools obtains access to student data for monitoring, how that data is analyzed, and how it is stored and secured.

> Bark — and the Bark for Schools initiative — is pioneering natural language processing systems to help keep kids and students safe in the digital age.

What is Bark?

Bark is an award-winning service that monitors texts, emails, YouTube, and 30+ apps and platforms for potential safety concerns. Created in collaboration with child psychologists, youth advisors, digital media experts, and law enforcement professionals, Bark delivers a research-backed solution for safeguarding children as technology changes how and where we communicate.

Launched in 2015, Bark looks for potentially harmful activities that may indicate online predators, adult content, sexting, cyberbullying, drug use, depression, suicidal ideation, and more. We created Bark to serve as an essential tool for parents raising kids in a digital age.

Bark for Schools

Bark's flagship product helps parents monitor their children's personal accounts. But in 2018, after the tragic Parkland shooting, we realized that we had the resources to provide the same service — in addition to others — to schools, and to do so in a way that keeps them from spending thousands of dollars each year to protect their students from the worst dangers of the online world. We developed Bark for Schools to give back to our communities and to help keep children safe — at absolutely no cost to them.

One of the fastest-growing companies in EdTech, we now help protect more than 2,000 school districts across the country — with dozens more joining every month.

G Suite + Office 365

Bark for Schools offers free monitoring of school-issued G Suite and Office 365 accounts, with specialized extensions for Chrome and Chromebook, to all public and private K-12 schools in the United States. We also provide a web filtering feature that allows administration to block and allow domains at the IP and DNS level.

> We developed Bark for Schools to give back to our communities and to help keep children safe – at absolutely no cost to them.

Data & Technology Overview

Data Control and Ownership

Schools and school districts are the absolute owners of all student data, and they also control the level of monitoring. Bark understands that there is no one-size-fits-all solution for school safety, and allowing each school to serve as the owner and controller of data and monitoring frees them to determine what's most appropriate for their particular institution.

We have implemented a number of features that make the Bark for Schools platform extremely easy to set up and use while maintaining a high standard of protection.

Chrome and Chromebook

Our monitoring service provides a specialized browser extension for Google Chrome, which is the inherent browser on Chromebooks. As Chromebooks are among the most popular devices issued to students in 1:1 districts, this extension allows Bark for Schools to monitor these browsers and devices with the same proficiency as with the broader G Suite account.

The benefits of a built-in monitoring extension for Chrome cannot be understated. While students can log out of G Suite or Office 365 accounts, Chromebook users are generally unable to circumvent our monitoring service because the browser is native to the device.

Other functionalities include:

- Account-level monitoring
- URLs
- Page titles
- Web searches
- Parent Portal-enabled

Clever + PowerSchool

Schools can use Clever to sync family contact information by downloading the Bark for Schools app. PowerSchool users can also connect to the Parent Portal by integrating Bark for Schools into their unified information system. These features greatly simplify the onboarding experience for schools, removing the need to manually input data and allowing administration to spend their time on more important things.

Our free service can be up and running within minutes, and once schools have synced their accounts, Bark for Schools will begin monitoring student activities for potential issues. Schools can also choose which apps and accounts to have covered by Bark for Schools. The login credentials for these platforms are not stored by Bark but are used to establish our access to their online activities.

Parent Portal

Students still use their school-issued accounts after classes have let out, and during those times there may not be a reviewer available to respond to urgent issues. The Parent Portal allows schools to enlist the support of parents and guardians to receive alerts about their children after hours and during breaks. Enabling the Parent Portal is a positive step towards more comprehensive monitoring to help keep students safe online.

Our Technology

Context is extremely important to properly classify a conversation. Without context, computers are unable to determine whether a specific chat message is a joke or something more serious. Modern machine-learning developments have allowed us to use cutting-edge, deep learning techniques to internalize that context, improving our ability to determine whether a piece of text is a potential issue.

Our underlying network is based on a recurrent neural network (RRN) trained on custom word embeddings to aid in understanding the variations of childspeak. We also use a multi-label approach, which allows us to use the triggers from one abuse type towards another. The combination of this approach trained on billions of data points has allowed us to create a perpetually improving natural language processing (NLP) system.

We've heavily invested in a robust Data Annotation team in the U.S. that gives us the capacity to properly label tens of thousands of pieces of data each day. This labeled data is then fed back into a data ingestion process that trains the algorithms nightly. Our investment in this process has given us the ability to stay current with the variations in language, as well as the flexibility to adapt on the fly as we encounter new topics and phrasings among student conversations.

Privacy, Security, Compliance

Security is one of our highest priorities at Bark, and we adhere to all of the standards required by FERPA, a federal law that protects the privacy of student education records. In an effort to provide maximum protection for students and their families, many states have adopted their own privacy legislation in addition to FERPA, and we meet or exceed each of these requirements, as well.

We secure all of our data within an encrypted database, including backups. The database is then purged after the data has been analyzed. Our server infrastructure is located in highly secured physical data centers that use a centralized bastion host, which is monitored to detect unwarranted access or activity. All web browser sessions use and require SSL encryption. And of course, no data is ever sold or given to any third party without consent.

> Our investment in this process has given us the ability to stay current with the variations in language.

Storage and Security

Web Browser Sessions

Web browser sessions are also encrypted and authenticated with SSL, meaning that all information moving between the web servers and browsers is kept completely private. No one can infiltrate these transfers without a unique cryptographic key issued by a Certificate Authority. Our certificate is issued by DigiCert, Inc., which uses the SHA-2 hashing algorithm to make it impossible for someone to modify or fake our certificate. Any such attempt triggers an error and prevents the attacker from making a secure connection.

Secure Databases

All data that is monitored by Bark for Schools is stored in an encrypted database. This includes field-level data like usernames and passwords, Google access tokens for Office 365, and other information that could be used inappropriately or without authorization. Some companies keep this kind of data stored in plain text, and when they're hacked it is easy to read and exploit that data. Bark for Schools has precluded this possibility, so a hacker would be unable to decrypt information stored in our database.

Physical Servers

Our physical servers are secured with a centralized bastion host, which essentially serves as a bridge between the Bark private network and the broader internet. The bastion host is configured for the specific purpose of withstanding cyberattacks. Relying on SSH for encryption and authentication, the bastion host also requires cryptographic keys to be accessed. Only Bark administrators have these keys, and we use centralized key management to grant or revoke access to them.

"The bastion host is configured for the specific purpose of withstanding cyberattacks.

Backups and Data Removal

The Bark for Schools databases are backed up every night and retained for a full week. They are also tested weekly to ensure that they can be restored. After seven days, the backups are purged entirely.

Data related to student activities is not included in the weekly purges, however. As noted previously, context is an essential component of our monitoring services. Our technology — as well as our team of human reviewers — requires a more protracted analysis in order to detect issues that develop over time. Accordingly, student activities are stored for 30 days, but at the request of a school or district, they may be removed in 15 days.

Amazon Web Services

Amazon Web Services (AWS) is a known and trusted partner in the cybersecurity industry, and they handle all of our database encryption needs. By working with such a reliable service provider, we are able to focus our efforts on Data Annotation, business development, and other necessary aspects of helping to keep children safe online.

Detection and Data Annotation

It's important to note that no monitoring service can promise complete and total detection of online issues. Adapting to the evolution of language is a perpetual challenge, and monitoring services are often limited by the closed nature of the platforms they are charged with monitoring. Snapchat and Instagram, for example, do not offer APIs that are conducive to third-party integration, which makes it difficult for monitors to detect every potential issue.

Furthermore, Bark allows schools and districts to exercise broad authority over which activities are monitored, as well as how sensitive our detection algorithms are to potential issues. Schools should be advised that adjusting the sensitivity levels has an effect on accuracy. Higher sensitivity may produce false positives, and lower sensitivity may produce false negatives.

That said, the years we've invested into refining our technology have produced an algorithm with a high rate of accuracy. In addition to the wealth of data collected for contextual analysis, this is in large part due to our dedicated Data Annotation team.

> Bark allows schools and districts to exercise broad authority over which activities are monitored.

Review and Escalation

The Data Annotation team consists of specialists trained to review and confirm instances of abuse that our algorithms detect. All team members have undergone a stringent background check to ensure their suitability for the highly sensitive nature of this work. They are also spread all across the U.S. so that they can be responsive to issues around the clock, regardless of the time zone, and they are all well-versed in the nuances of contemporary American slang. By confirming or overruling alerts to abuse, the team contributes to the daily improvement of our algorithms.

The capabilities of our algorithms are such that 99.96% of messages are not seen by human eyes. Reviewers on the Data Annotation team only encounter anonymized student data, and only in the event of a severe escalation are the review team leaders able to see more specific information. By studying communications at the human level, they are able to make more informed decisions and ultimately coordinate the appropriate response to prevent a tragedy from unfolding.

Bark works with law enforcement at the federal, state, and local levels, and our team escalates cases of potential grooming, exploitation, drug trafficking, or imminent risk to a student's life to the appropriate authorities. We also work with the National Center for Missing and Exploited Children (NCMEC) in cases involving sexual exploitation.

After issuing a severe alert, we keep track of it for 24 hours. If schools don't review the issue within that time, we email them. The Data Annotation team leaders send these emails after judging that the issue is serious enough to merit the school's immediate involvement. In escalated situations, automatic email and text alerts notify the relevant parties so they can be aware of the situation and the individuals involved.

We highly recommend that schools have a **designated person** to field alerts from Bark. We also recommend that schools have **at least two contacts**, as well as a **mobile phone number** listed for immediate communications with Bark. Severe issues — like a school shooting threat — are often time sensitive, so it's urgent that a contact is available. Issues may also occur after school hours and during holidays, when many students are still using their school-issued accounts. Schools can allow parents to receive alerts during these times by **enabling the Parent Portal**.

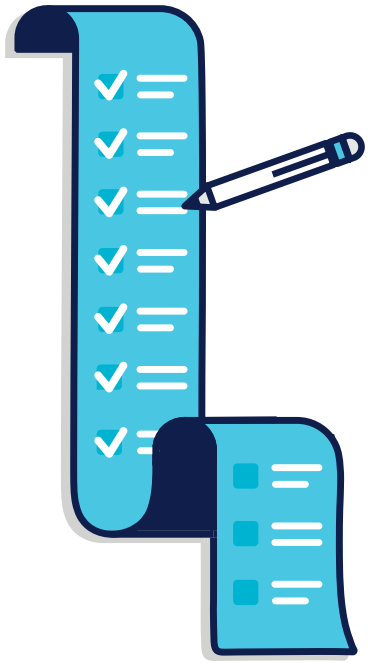

Severity and Type

Activities that Bark's algorithms label as Abusive are categorized by "Severity" and "Type." These two measures are weighted so they can be scored on the same scale. If the aggregated severity level is high enough then it will be automatically escalated to the Data Annotation team. Alerts are also given a Confidence Score, which is the computer's level of certainty that the issue is legitimate. If the Confidence Score is high enough, the alert will be auto-sent, notifying Reviewers directly through the Bark for Schools platform.

Severity is a measure of the seriousness of a potential issue, as judged by our technology's contextual analysis — the higher the number, the more severe the alert. For example, a student sharing plans to bring a gun to school is more severe than a student wishing a classmate would die.

Type describes what kind of issue has been detected, like "Cyberbullying" or "Suicide / self-harm." Each type of abuse has a different threshold for escalation, and between the provided examples, suicidal ideation would be rated higher than an instance of cyberbullying.

Abuse Types include:

- Cyberbullying
- Sexual content
- Depression
- Suicide / self-harm
- Drug-related
- Violence
- Hate Speech
- Grooming and Sextortion
- Other (including profanity, drug use, alcohol, etc.)

Reports and Analytics

Depending on the amount of activity, reports can be sent either as a lump sum of aggregated data, or by directing the requested data upstream to be generated into reports on a rolling basis. For example, if a fight breaks out at school, the administration may want to investigate the role that cyberbullying played in the event — the number of total instances, what time of day or night they occured, and/or which accounts and platforms were used — so they can develop a more informed response.

The school can request a report on all incidences of cyberbullying (within 30 days), and Bark will automatically compile that data for their review. They can also cross-reference other abuse types, like Depression or Hate Speech, if their investigation leads them to do so. These reports can be a valuable tool for schools to understand the dynamics of their student bodies.

Remember: schools and districts own all student data. Bark for Schools monitors that data on their behalf to help keep students safe.

State and Federal Compliance

Privacy and data security are at the top of people's minds when it comes to their online activities. In light of hacks and data breaches of increasing severity, individuals and organizations alike want to ensure that they are protecting both themselves and those who depend on them.

Signing up for Bark for Schools helps keep students safe online, and the data we collect is secured with state-of-the-art technologies and best practices. Bark for Schools can also help schools meet certain state and federal compliance regulations.

Children's Internet Protection Act (CIPA)

Protects minors from obscene or otherwise harmful online content at school.

Requires adoption of an internet safety policy that:

- Blocks or filters inappropriate content
- Monitors students' online activities
- Provides education about appropriate online behavior

How Bark for Schools Can Help

Bark for Schools provides free monitoring services for all K-12 accounts, including email, documents, and cloud storage solutions offered by G Suite and Office 365. We also filter web domains at the IP and DNS levels.

Children's Online Privacy Protection Act (COPPA)

Prohibits deceptive online practices for the collection, use, or disclosure of personal information of children under 13. Site operators must provide parental notice and consent, review processes, confidentiality, data retention and deletion statements, and more.

How Bark for Schools Can Help

Bark for Schools does not collect data from children under 13 without parental consent. Data is also automatically deleted after 30 days, and upon request, user data may be deleted within 15 days.

Federal Educational Rights and Privacy Act (FERPA)

Provides rights to parents and eligible students over education records, including the right to review and correct the student's records. Schools must generally have written consent to release information from a student's record.

How Bark for Schools Can Help

Bark for Schools does not disclose information to any third party without prior written consent unless pursuant to court or administrative order. Bark for Schools is also considered a "School Official" with legitimate educational interest under FERPA and may share without consent to protect the student body.

California Assembly Bill No. 1584

Authorizes educational agencies to contract with third-party providers under specific requirements. Applies to providers of electronic services for:

- Digital storage, management, and retrieval of student records
- Educational software that allows third parties to access, store, and use student records

How Bark for Schools Can Help

Bark for Schools is an at-will service provider and is therefore not subject to third-party contract requirements. But because Bark for Schools complies with COPPA and FERPA, it also satisfies several tenets of AB 1584.

Bark for Schools Privacy and Security Checklist

- ✔ Schools and districts own all user data
- ✔ Schools control the level and degree of monitoring
- ✔ Data stored in encrypted database
- ✔ Web browser sessions use SSL encryption
- ✔ Physical servers secured with centralized bastion host
- ✔ Data purged within 30 days of analysis
- ✔ No data sold or given to third parties without consent
- ✔ Consent may be withdrawn at any time

Visit our website at **bark.us/schools** for more information about Bark for Schools.

bark FOR SCHOOLS