Private Compute Core Architecture

Eugenio Marchiori, Sarah de Haas, Sergey Volnov,
Ronnie Falcon, Roxanne Pinto, and Marco Zamarato

Google

September 23, 2022

1 Introduction

Android's Private Compute Core (PCC) is a secure, isolated environment within the operating system, that maintains separation from apps while enabling users and developers to maintain control over their data [1]. It is backed by open-source code in the Android Framework introduced in Android 12 [2]. PCC allows features to communicate with a server to receive model updates and contribute to global model training through Private Compute Services (PCS), the core of which has been open sourced[1].

PCC is part of the OS, and by virtue of being isolated, constrained, and trusted, it can host sophisticated ML features. The hosted features themselves, running inside PCC, can be closed source and updatable. In this way, PCC enables machine learning features to process ambient and OS-level data and improve over time, while restricting the availability of information about individual users to servers or apps.

With each new Android release we'll add more features to the Private Compute Core. Some examples of features that now operate within PCC include:

- Live Caption, which adds captions to any media using Google's on-device speech recognition

- Now Playing, which recognizes music playing nearby and displays the song title and artist name on the device's lock screen

- Smart Reply, which suggests relevant responses based on the conversation you're having in messaging apps

2 System architecture

2.1 Data handled/managed by PCC

In this document, "ambient and OS-level data" refers to three main categories:

- **Raw data:** this is usually from the device sensors (camera, microphone, etc) or generated by the OS for its normal operation (such as screen content). This data is shared by the operating system with PCC so that it can manage feature access and its usage. This is referred to as "ambient and OS-level data" or just "data" throughout this document.

- **Derived data:** this is data generated from analysis or inferences based on ambient and OS-level data. Derived data can in some cases be allowed to leave the PCC environment, subject to technically enforced policy.

---

[1]Details of what has been open sourced can be reviewed at Private Compute Services. Google Open Source Project on GitHub. Retrieved Jul 2022 from https://github.com/google/private-compute-services.

- **Metadata from system operation** this data is considered non-user identifiable if the user retains k-anonymity[3] (where k is feature and context dependent), but some can be identifiable and thus subject to stricter precautionary measures.

The following is a non-exhaustive categorization of collected ambient and OS-level data in PCC:

CATEGORY	TYPE	SOURCE	METHOD OF ACCESS
OS-level	audio	framework audio (e.g. song currently playing)	via audio APIs
	image	screenshots (i.e. when changing activity)	via Content Suggestions API
	text	screen capture	via Content Capture API
		notification content	via NLS[a] and NAS[b]
	structured	app data (indexed)	via AppSearch API
		app data (pushed)	via ContentCapture API
		contacts	via Contacts API
		app launches	via UsageStatsManager
		shortcuts	via ShortcutManager
ambient	audio	background microphone	via audio APIs
	image	background camera	via camera APIs
	location	GPS and other signals	via Location Provider

[a]NotificationListenerService
[b]NotificationAssistantService

Table 1: Types and sources of data used in PCC

2.2 Overview diagram

The current high-level architecture of Private Compute Core is pictured in Figure 1. PCC exists in its own isolated virtual sandbox (following a verifiably private design), with features allowed to run computations inside that sandbox on ambient or OS-level data generated by the user. The results are surfaced to the user via the UI operated by the trusted OS, or through access-controlled (by permissions or other mechanisms) open-source framework APIs as needed[2].

The inferences from ambient and OS-level data are retained inside Android System Intelligence, within Private Compute Core. All communications to external servers is done via Private Compute Services (PCS) privacy-preserving open-source technologies.
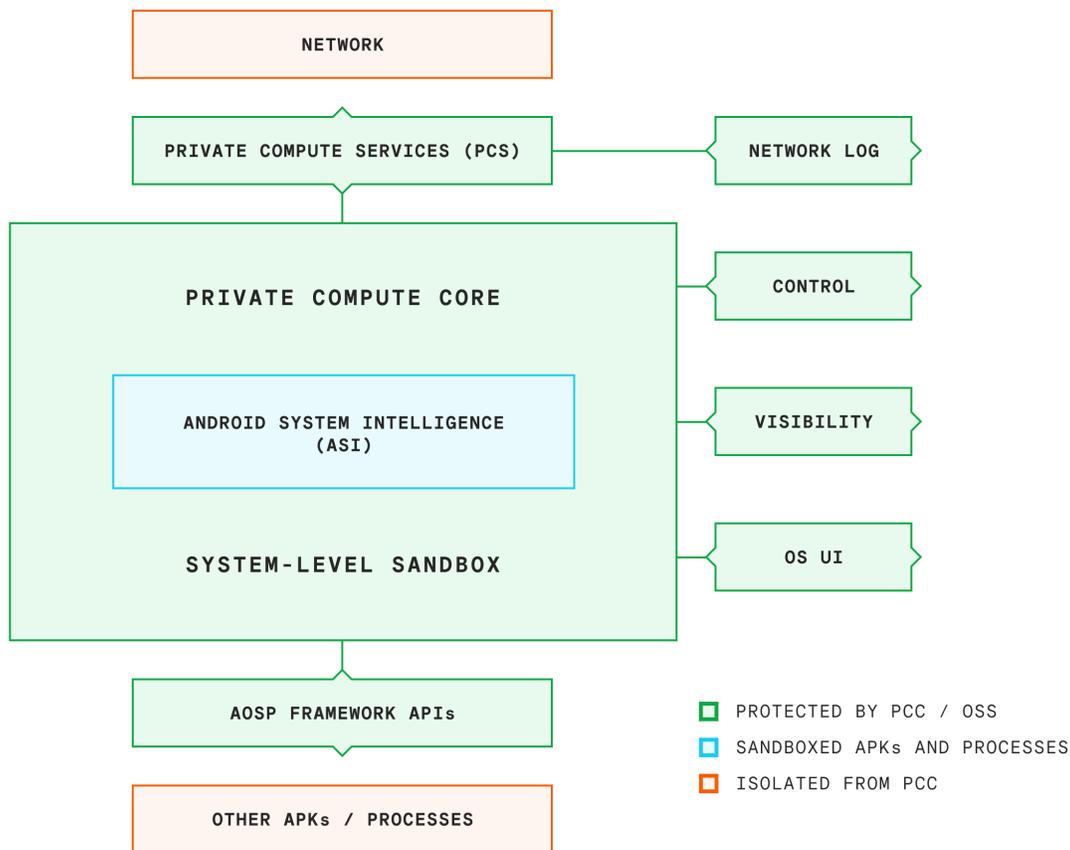

Figure 1: Conceptual architecture diagram of Private Compute Core

2.2.1 Android System Intelligence

Android System Intelligence (ASI) is a Google-proprietary package containing system-level features like Live Caption[4], Now Playing[5], and others. It can only be accessed on compatible devices, and the available features may vary.

---

[2]Minimal logging uses Android Framework APIs and is limited to non-identifiable data such as performance and stability metrics.

ASI restricts usage of the data through mechanisms like the `<allow-association>` system (Table 2), or the absence of `android.permission.INTERNET`[3] permission (through the system permission system). It solely relies on PCS (described in detail below) and the Android framework for any information egress, both of which are open-source.

2.2.2 Private Compute Services

Starting in Android 12, features within PCC no longer have `android.permission.INTERNET` permission, to allow claims about what data is shared with the network to be externally verifiable. To retain functionality without internet capability and increase the transparency of how features use network access, Private Compute Services (PCS) was created.

Private Compute Services (PCS) is an APK with open-source[4], code, which provides a narrow set of structured APIs purposefully designed for PCC components. Currently, PCS supports:

1. **Federated learning[6] and federated analytics[7],** enabling privacy-preserving machine learning and analytics without centralized data collection. The underlying techniques involve pushing a computation graph (e.g. machine learning model) to the device, computing on the locally stored data, and sending only the computation results back. The results from many devices are aggregated together, and used to improve the device features and user experience. Each individual device's results are protected from being seen by the orchestrating server through the use of the Secure Aggregation[8] multi-party computation protocol, ensuring that only aggregates over many (e.g. thousands) of devices are made available to servers and model/feature developers.

2. **Private Information Retrieval (PIR)[9],** enabling privacy-preserving lookups of information from the cloud. The PIR protocol relies on a client-side implementation of homomorphic encryption to retrieve resources from a server-hosted database, while cryptographically ensuring that the server is unable to observe which resource was downloaded.

3. **An HTTPS download-only transport,** which limits downloads to pre-specified global ML models.

2.3 Isolation mechanisms

2.3.1 IPC bind restrictions

PCC sets constraints on data sharing both for off-device services and those working on-device. On the client, technical limits are set on the other processes in the operating system (including apps) that PCC features can communicate with.

This is called the Inter Process Communication (IPC) allow-association mechanism, and it defines the exhaustive list (allow-list) of packages that a PCC package can establish an IPC channel with (outside the existing Android API surface), or vice versa. The recommended list of packages in the allow-list is defined in the CDD (Compatibility Definition Document) under ContentCapture and AppSearch.

As a result, only the system can establish communications between PCC and other services, which allows for inspectable API calls with any process not in the allow-list.

2.3.2 Process isolation

Some instantiations of PCC are not applicable at the package level, but instead only affect a single feature. For these, the system can use Android's isolated processes[5][10], in conjunction with other means (e.g. permission management), which are designed to ensure that the data remains separate from the application

---

[3]https://developer.android.com/reference/android/Manifest.permission#INTERNET
[4]PCS code is continuously published to https://github.com/google/private-compute-services/
[5]https://developer.android.com/guide/topics/manifest/service-element#isolated

shipping the code and all APIs are open-source and inspectable.

Through this mechanism, it's possible to build features which follow the PCC implementation pattern, but whose code might be provided by an APK that fulfills other roles. Some versions of this also force these processes to be ephemeral (starting new ones regularly to replace the old ones).

2.4 Permissions and controls

2.4.1 Android permissions

Data going in and out of PCC features does so via Android Framework APIs. By extension, this means that such data flows are governed by the Android permission framework[6]. Given that most sensitive permissions can only be pre-granted by the OS (via explicit allow lists or roles), this can't be modified after the fact, except on an operating system update. Anyone is able to analyze requested permissions by PCC components on a device at install time (by looking at the manifest) and runtime (via Android Debug Bridge[7]).

2.4.2 Android platform controls

PCC features depend on APIs defined in the platform. Other implementations of features in the Android ecosystem that use screen content or app search data are subject to these controls, as they are required to pass CTS (Compatibility Test Suite) and comply with Android's CDD (Compatibility Definition Document).

Android System Intelligence functions as a non-replaceable system component. Although it has its own separate update process, Android will check signatures upon each update making sure that it cannot be easily replaced by implementations from a different source.

2.4.3 User controls

PCC features allow data to leave the designated sandboxes under limited circumstances. When possible the controls are implemented as part of the platform and not managed by the PCC feature itself.

PCC implementations offer additional ways for users to control if and how features use their data, and management of the data retained by those features. These include:

1. In Android OS, "Settings→Privacy→Personalize using app data". This setting is part of AOSP and restricts screen data, app search data, and most other app-linked data from being shared with PCC features.

2. "Clear data" controls, offered by the features themselves. These are usually not open-source, but allow for granular control in many cases, including deleting all data or only for a period of time.

3. Android runtime permissions, controlled by AOSP, which the user can revoke.

4. Android sensor toggles, which allow the user to disable the camera or microphone, and prevent the data from these sensors from being used across the OS including by PCC features.

2.4.4 App developer controls

Following the Android Platform Security Model[10] developers maintain control of data generated by their own applications.

An example of a control is provided by `FLAG_SECURE`[8]. This is a mechanism, defined in Android, that allows apps to opt-out of operating system functionalities (in this case screenshots or being viewed on non-secure displays). By extension, this also applies to data processing in PCC from activities included under this setting. In some cases, more granular controls are provided on a per-feature basis.

---

[6]https://developer.android.com/guide/topics/permissions/overview
[7]https://developer.android.com/studio/command-line/adb
[8]https://developer.android.com/reference/android/view/WindowManager.LayoutParams#FLAG_SECURE

2.5 Surfaces

PCC features might interact with surfaces that are outside of the PCC environment. To achieve this, we define two types of surfaces - Public Framework APIs and Operating System User Interfaces (which also use APIs):

2.5.1 Public Framework APIs

In some cases we expose PCC features through Framework APIs (for example, a speech recognizer[9] which ensures speech data is protected). For cases like these, we implement the following:

1. *Framework proxies:* to make sure that feature updates still follow the same rules, all API communications with PCC features are proxied by the Android platform. This means that the platform can enforce rules on how they are used (e.g. rate limiting) and what data is passed on to other apps (e.g. no side channels).

2. *Data limitations:* data can only be used if it will not result in user identification or escalation of permission by the caller app.

2.5.2 Operating System User Interfaces

As a general rule, exposing PCC data in OS surfaces is treated differently than when PCC data is exposed to apps. Exposure in the OS is considered trusted and necessary for technical and practical reasons, such as needing to surface a sharing permission request to the user. However, when data is exposed to users in a UI outside PCC there are three options to consider:

1. *Appearance in an OEM-rendered system UI.* These UIs usually include containers like the app launcher or notifications panel, and the UI shows data from PCC via System APIs with specific purposes.

2. *Appearance in AOSP-rendered system UI.* This includes a growing set of surfaces like the "Privacy dashboard" where data can be safely shown to the user.

3. *Appearance in delegated UI.* This is a mechanism where the rendering is done by the PCC feature but shown in a different app, while restricting the target app from observing the content or user actions via the isolation and controls mechanisms previously described.

2.6 Data sources

There are several existing Android APIs that are used as data sources for PCC features. In general, users and apps can opt-out of data being shared over these APIs. Two examples of these data source APIs, for which the requirements are described in the Android Compatibility Definition Document (CDD), are described below.

2.6.1 Content Capture API

Content Capture API is a mechanism introduced in Android 11 to allow the OS and PCC features to use screen data (under user and application control). Usage of this data is very limited even within PCC, and never shared across apps without specific user actions. Features such as Smart Reply, described in detail later, rely on this API.

In Android, the standard user interface tools are instrumented so that data is automatically formatted and made available to the API when rendered on the screen. The API can also include other types of

---

[9]https://developer.android.com/reference/android/speech/SpeechRecognizer.html#createOnDeviceSpeechRecognizer(android.content.Context)

structured app data, specifically passed by the apps to PCC (i.e. share-data)[10].

Third-party applications can always opt-out of their data being passed through the API to PCC features, and this is automatically enforced by the operating system. Applications can also request that all instances of a specific piece of data (e.g. a message that was deleted by the user) be deleted by specifying its unique "locus id"[11] generated by the API.

2.6.2 AppSearch

AppSearch is a general purpose indexing and searching mechanism, provided by the Android OS to all app developers[12]. This is designed to run optimally on-device and fulfill most applications needs. It was introduced in Android 12, with backwards compatibility via Android Jetpack[13].

AppSearch also allows data to be shared with different services, and it is shared with PCC on devices where it is available, but application developers can opt-out of this.

3 Implementation examples

Following the PCC architecture previously described, Android System Intelligence (ASI) sandboxing is achieved by the system rules, for example `<allow-association>` system (Table 2), or the absence of `android.permission.INTERNET`[14] permission (through the permission system). It solely relies on PCS and the Android framework (via different APIs) for any information egress.

```
<allow-association target="com.google.android.as"allowed="com.android.bluetooth"/>

<allow-association target="com.google.android.as" allowed="com.android.providers.contacts" />

<allow-association target="com.google.android.as" allowed="com.android.providers.media"/>

<allow-association target="com.google.android.as" allowed="com.android.providers.telephony"/>

<allow-association target="com.google.android.as" allowed="com.android.systemui"/>

allow-association target="com.google.android.as" allowed="com.google.android.providers.media.module"/>

// Private Compute Services

allow-association target="com.google.android.as" allowed="com.google.android.as.oss"/>
```

Table 2: Allow-association configuration for ASI

The `<allow-association>` mechanism blocks any incoming and outgoing inter-process requests through primitives like `bindService()`, using content providers or content resolvers, broadcasting intents and others.

In addition to the `<allow-association>` mechanism Android System Intelligence follows the Android permission model. The APK publicly requests the required permissions (Table 3) in its manifest, and are typically pre-granted by the OS (so ASI must pass both the association and any traditional permission checks). In general, it is considered that data access permissions (e.g. `android.permission.ACCESS_WIFI_STATE`) are safe to be granted to the APK given the technical limitations for data sharing imposed by PCC.

---

[10]https://developer.android.com/reference/android/view/contentcapture/ContentCaptureManager#shareData(android.view.contentcapture.DataShareRequest,%20java.util.concurrent.Executor,%20android.view.contentcapture.DataShareWriteAdapter)

[11]https://developer.android.com/reference/android/content/LocusId

[12]https://developer.android.com/guide/topics/search/appsearch

[13]The backwards compatibility implementation via Android Jetpack means that prior to Android 12, data is not shared outside of the client app.

[14]https://developer.android.com/reference/android/Manifest.permission#INTERNET

On Pixel devices ASI APK fulfills the following roles[15], defined by the OS[16]:

System UI Intelligence (SYSTEM_UI_INTELLIGENCE)

System Ambient Audio Intelligence (SYSTEM_AMBIENT_AUDIO_INTELLIGENCE)

System Audio Intelligence (SYSTEM_AUDIO_INTELLIGENCE)

System Notification Intelligence (SYSTEM_NOTIFICATION_INTELLIGENCE)

System Text Intelligence (SYSTEM_TEXT_INTELLIGENCE)

System Visual Intelligence (SYSTEM_VISUAL_INTELLIGENCE)

```
android.permission.ACCESS_BACKGROUND_LOCATION
android.permission.ACCESS_COARSE_LOCATION
android.permission.ACCESS_FINE_LOCATION
android.permission.ACCESS_NETWORK_STATE
android.permission.ACCESS_SHORTCUTS
android.permission.ACCESS_WIFI_STATE
android.permission.BLUETOOTH_ADMIN
android.permission.BLUETOOTH_CONNECT
android.permission.BLUETOOTH_SCAN
android.permission.CAMERA
android.permission.CAPTURE_AUDIO_HOTWORD
android.permission.CAPTURE_AUDIO_OUTPUT
android.permission.CAPTURE_MEDIA_OUTPUT
android.permission.CAPTURE_VOICE_COMMUNICATION_OUTPUT
android.permission.CONTROL_INCALL_EXPERIENCE
android.permission.EXEMPT_FROM_AUDIO_RECORD_RESTRICTIONS
android.permission.FOREGROUND_SERVICE
android.permission.MANAGE_APP_PREDICTIONS
android.permission.MANAGE_MUSIC_RECOGNITION
android.permission.MANAGE_SEARCH_UI
android.permission.MANAGE_SMARTSPACE
android.permission.MANAGE_SOUND_TRIGGER
android.permission.MANAGE_UI_TRANSLATION
android.permission.MODIFY_AUDIO_ROUTING
android.permission.MODIFY_AUDIO_SETTINGS
android.permission.MODIFY_PHONE_STATE
android.permission.OBSERVE_SENSOR_PRIVACY
android.permission.PACKAGE_USAGE_STATS
android.permission.QUERY_ALL_PACKAGES
android.permission.READ_CALL_LOG
android.permission.READ_CONTACTS
android.permission.READ_DEVICE_CONFIG
android.permission.READ_EXTERNAL_STORAGE
android.permission.READ_GLOBAL_APP_SEARCH_DATA
android.permission.READ_OEM_UNLOCK_STATE
android.permission.READ_PEOPLE_DATA
android.permission.READ_PHONE_STATE
android.permission.READ_SMS
android.permission.RECEIVE_BOOT_COMPLETED
android.permission.RECORD_AUDIO
android.permission.REQUEST_NOTIFICATION_ASSISTANT_SERVICE
android.permission.START_ACTIVITIES_FROM_BACKGROUND
android.permission.SUBSTITUTE_NOTIFICATION_APP_NAME
android.permission.SYSTEM_APPLICATION_OVERLAY
android.permission.SYSTEM_CAMERA
android.permission.UNLIMITED_SHORTCUTS_API_CALLS
android.permission.UPDATE_DEVICE_STATS
android.permission.VIBRATE
android.permission.WAKE_LOCK
android.permission.WRITE_SECURE_SETTINGS
com.android.alarm.permission.SET_ALARM
com.google.android.ambientindication.permission.AMBIENT_INDICATION
com.google.android.apps.nexuslauncher.permission.HOTSEAT_EDU
com.google.android.setupwizard.SETUP_COMPAT_SERVICE
```

Table 3: Permissions requested by ASI APK, at the time of writing

---

[15]https://developer.android.com/reference/android/app/role/RoleManager
[16]https://cs.android.com/android/platform/superproject/+/master:packages/modules/Permission/PermissionController/res/xml/roles.xml

ASI is covered by the CDD for `ContentCaptureService`, `AppSearch` Services[17] and Android Intelligence roles[18] as it has access to screen and app search content, and other data. This means that ASI follows CDD guidance such as:

> 9.8.6 [C-SR] Are STRONGLY RECOMMENDED NOT to request the INTERNET permission.
>
> 9.8.6 [C-SR] Are STRONGLY RECOMMENDED to only access the internet through structured APIs backed by publicly available open-source implementations.`
>
> 9.8.6 [C-SR] Are STRONGLY RECOMMENDED to keep the services separate from other system components(e.g. not binding the service or sharing process IDs) except for the following: Telephony, Contacts, System UI, and Media

These requirements are enforced (i.e. turned into MUST instead of STRONGLY RECOMMENDED) for packages holding the `SYSTEM_*_INTELLIGENCE` roles (mentioned above).

Below are some example features provided by Android System Intelligence as implemented in Pixel and potentially other devices for Android 12.

---

[17]https://source.android.com/docs/compatibility/12/android-12-cdd#986_content_capture_and_app_search
[18]https://source.android.com/compatibility/12/android-12-cdd#91_permissions: [C-4-1...C-4-3]

3.1 Smart Reply

Smart Reply (part of Smart Input[11]) suggests quick responses on the keyboard based on both previous and current screen content. To do this, ASI extracts relevant entities from different apps such as addresses, names, and others, and then offers them to the user as suggestions. These suggestions show up in compatible keyboards, but leverage privacy preserving approaches to avoid revealing the content to the keyboard systems.
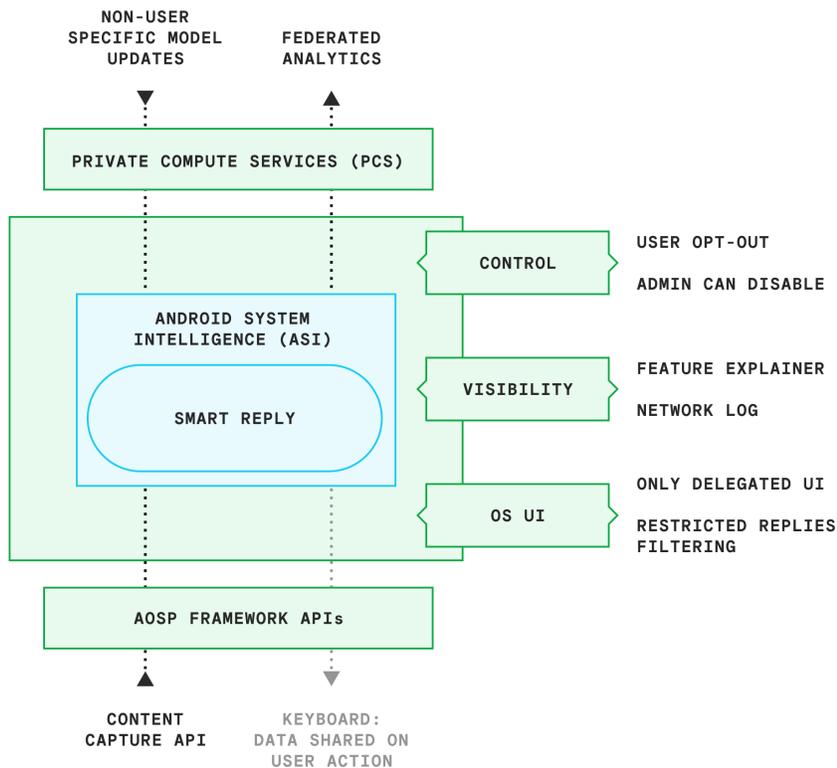

Figure 2: Smart Reply Implementation

This feature is implemented as follows:

- Uses Content Capture API as a data source, which is an Android Framework API with clear control in the framework. This ensures limits on access by the feature.

- Users (via Settings) and app-developers (via `FLAG_SECURE` or overriding the API) can opt-out of screen data being used by this feature.

- Device administrators can disable screen capture for enterprise devices by using the `POLICY_DISABLE_SCREEN_CAPTURE`[19] policy or through Android Management[20] cloud API. This ensures that admins retain control of what data is accessible to features and services.

- The functionality is explained[11] and can be disabled by the user in the UI.

---

[19]https://developer.android.com/reference/android/app/admin/DevicePolicyManager#POLICY_DISABLE_SCREEN_CAPTURE

[20]https://developers.google.com/android/management/reference/rest/v1/enterprises.policies

- Users can specifically allow data to leave PCC

  - To display the candidate responses to the user in a keyboard which supports it such as Gboard, a delegated UI via a specific Android Framework API is used, which means that no data leaves the PCC boundary at rendering time.

  - Candidate filtering by input is disabled after a series of keystrokes, due to the keyboard being able to obtain how many candidate responses are displayed. This prevents attempts to guess the content of the responses prior to the user tapping on the chip.

  - Data is only held ephemerally in PCC (kept in memory with short TTLs[21], thus suggestions are based on the data observed recently).

  - Data is only obtained from apps that PCC understands how to infer data from (by supplying an allowlist to the system).

- Uses Private Compute Services' OSS APIs for all network access

  - ML models are non-user specific.

  - Analytics is performed through federated analytics with secure aggregation, with additional limitations to reduce risk (e.g. restricting the analysis to only popular apps, limited TTLs of data on-device, etc). Data semantics being used for federated analytics is documented in the open-source code in a form of policy that is enforced in the close-source code in PCC.

---

[21]TTL or Time-To-Live refers to the time that data is kept before being deleted.

3.2 Live Caption

This feature provides captions for any content playing on the phone, using best-in-class on-device speech recognition. To do this, it processes all audio playing on the Android device and displays the transcript as an overlay which is an AOSP-rendered UI (data is not accessible to apps).
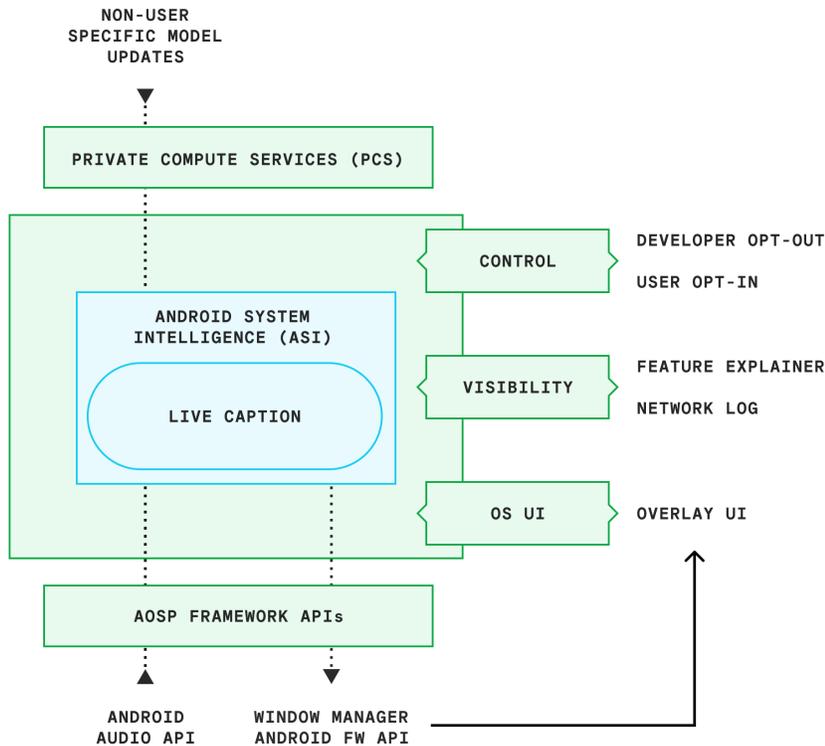

Figure 3: Live Caption Implementation

This feature is implemented as follows:

- Uses Android Audio APIs as a data source, which is an Android Framework API with clear control in the framework, which ensures limits on access by the feature.

  – App-developers (via `ALLOW_CAPTURE_BY_NONE`) can opt-out of data being used by this feature.

- The feature is turned on by the user and shows a clear UI when on. The Live Caption settings screen details how the feature works and what data is used.

- Data does not leave PCC and is only rendered in a system surface

  – To display the transcript the feature uses a Window Manager Android Framework API to create an overlay.

  – Data is only held ephemerally in PCC (kept in memory with short TTLs).

- Uses Private Compute Services' OSS APIs for all network access

  – Model updates are non-user specific.

3.3 Screen attention

When enabled, this feature prevents screen deactivation while the user is still looking at their phone. At the time screen dimming is scheduled to happen, if a face is detected in the front-facing camera's view, then the dimming is postponed.
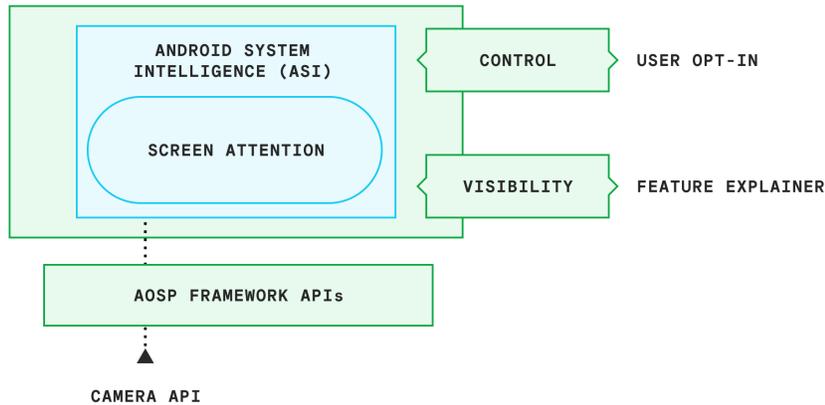

Figure 4: Screen Attention Implementation

This feature is implemented as follows:

- Uses standard Camera APIs as a data source, which is an Android Framework API with clear control in the framework. This ensures limits on access by the feature.

- The user can turn on this feature in a UI[22] where it is clearly explained. The feature is not enabled without this explicit opt-in.

- Data does not leave PCC and is only processed within PCC and the OS through the `AttentionManagerService` Framework API. Data is only held ephemerally (kept in memory with short TTLs).

- Does not use any of the network capabilities, as the models are only updated via the APK.

---

[22]https://support.google.com/pixelphone/answer/6111557

4   Conclusion

Private Compute Core, as defined in Android 12, has introduced mechanisms that allow for the creation of high-value features that use ambient and OS-level data while following Android guidelines and privacy-preserving methodologies. Thanks to Private Compute Services, these features can still use cloud services in limited and privacy-preserving ways, through OSS that can be reviewed by external experts. Android System Intelligence, as available on all Pixel phones and other Android devices, represents the main vehicle for PCC features. As PCC grows, we expect to expand the space and include many new features, following and improving upon the mechanisms outlined in this paper.

5   Acknowledgements

References

[1]   Suzanne Frey. *Introducing Android's Private Compute Services.* https://security.googleblog.com/2021/09/introducing-androids-private-compute.html. Google Security Blog. Sept. 2021.

[2]   Sameer Samat. *Android 12 Beta: Designed for you.* https://blog.google/products/android/android-12-beta/. The Keyword Blog. May 2021.

[3]   Latanya Sweeney. "K-Anonymity: A Model for Protecting Privacy". In: *International Journal of Uncertainty, Fuzziness and Knowledge-Based Systems* 10.05 (2002). DOI: https://doi.org/10.1142/S0218488502001648, pp. 557–570.

[4]   Michelle Tadmor-Ramanovich and Nadav Bar. *On-Device Captioning with Live Caption.* https://ai.googleblog.com/2019/10/on-device-captioning-with-live-caption.html. Google AI Blog. Oct. 2019.

[5]   Beat Gfeller et al. *Now Playing: Continuous low-power music recognition.* https://research.google/pubs/pub46522. NIPS 2017 Workshop: Machine Learning on the Phone. Oct. 2017.

[6]   Brendan McMahan and Daniel Ramage. *Federated Learning: Collaborative Machine Learning without Centralized Training Data.* https://ai.googleblog.com/2017/04/federated-learning-collaborative.html. Google AI Blog. Apr. 2017.

[7]   Daniel Ramage and Stefano Mazzocchi. *Federated Analytics: Collaborative Data Science without Data Collection.* https://ai.googleblog.com/2020/05/federated-analytics-collaborative-data.html. Google AI Blog. May 2020.

[8]   K. A. Bonawitz et al. *Practical Secure Aggregation for Federated Learning on User-Held Data.* https://research.google/pubs/pub45808/. NIPS Workshop on Private Multi-Party Machine Learning. 2016.

[9]   Benny Chor et al. "Private information retrieval". In: *In Proceedings of IEEE 36th Annual Foundations of Computer Science* (Oct. 1995). DOI: https://doi.org/10.1109/SFCS.1995.492461.

[10]   René Mayrhofer et al. *The Android Platform Security Model.* https://arxiv.org/pdf/1904.05572.pdf. 2020.

[11]   Yang Lu, Angana Ghosh, and Xu Liu. *Privacy-Preserving Smart Input with Gboard.* https://security.googleblog.com/2020/10/privacy-preserving-smart-input-with.html. Google Security Blog. May 2020.