

Uttu

Customer Data & Analytic Platform

v1.0a

October 2018

Protocol One

Who might be interested in this document

Data presented in this document is relevant to a wide range of video game, advertising and analytics market participants:

- analysts and marketers;
- publishers and game developers;
- publishing platforms;
- technical experts;

This document describes the technical and procedural aspects of the Uttu component of the Protocol One ecosystem. Description of the Ecosystem can be found within the [Protocol One Whitepaper](#) document.

Introduction

Uttu is an open analytical and marketing platform for collecting and analyzing data, supporting end-to-end users attribution with data enrichment from external sources and forecast services. Uttu allows to automate marketing and to build flexible scenarios for active interaction with the audience.

Uttu is supplied with open source code under the Apache 2.0 license. The platform can be installed free of charge, and used either on the owner's hardware, or within a cloud.

The Protocol One ecosystem provides a "Glutton Uttu" cloud solution of its own, using the SAAS model. Additional information concerning differences in installations can be found in the "[Glutton Uttu](#)" section.

Uttu is capable of:

- Giving instant answers to tough questions to users lacking technical education, based on raw data and in real time. No sampling and delays in incoming data.
- Collecting and processing data for cross-domain, cross-channel tracking on the web, desktop, mobile applications and games.
- Building any kind of slices based on selected data. There are no restrictions on the amount of event and metric parameters in Uttu.
- Automatically enriching data from external sources and storing a 360° profile in one place.
- Easily integrating data with one's own CRM or e-commerce solution.
- Collecting full-cycle data from web pages to in-game user behavior, so that the analyst can directly connect the channels the users came from to their behavior.
- Segmenting visitors to determine their interests, preferences and behavioral patterns; gathering users into cohorts.
- Calculating the optimal way to interact with the user and sending data in real-time to browser, where it can be used to make decisions.
- Using machine data analysis and predictors to find patterns in a large amount of data.

The [Open Game Data Exchange](#) format is used to facilitate data exchange pertaining to games and the audience, and interaction between the participants.

Roadmap

We involve as many professionals and companies working with data as possible to make Uttu better and easier to use. At this moment Uttu is under active development, preparing for a public launch. The roadmap for the next 4 iterations of Uttu is outlined below.

Version	Date	Key tasks
1.0	November 2018	Uttu specifications, implementation of data collection and normalization, basic reports and BI dashboards.
1.0b	March 2019	First version of the product. Integration with GameNet, Super.com. Basic set of predictive services, implementation of the scenario construction set. Basic integration of CPA networks specified in this document.
1.0rc	July 2019	Integration of advertising networks, first implementation of the Open RTM connector. Implementation of built-in tools for BI analytics. Preparation for release.
1.0r	November 2019	Release and the start of commercial exploitation.

Introduction	2
Roadmap	3
Starting work	5
What happens to the data?	5
Session	6
Segments	6
Structure	7
SDK	8
Bulk Import API	9
Data Hub	9
Analytics Tools	9
Scenario Construction Set	10
Toolbar	11
Multichannel sequences	12
Working with raw data	13
SQL support	13
Predictive services	14
Ad networks	15
Architecture	16
Technology stack	16
Event handling	16
Suggested hardware	16
Scalability	17
Data storage	17
Data storage on disk	17
Distributed query processing	18
Data protection	18
Glutton Uttu	18
Disclaimer	20
Appendix 1. Default settings	21

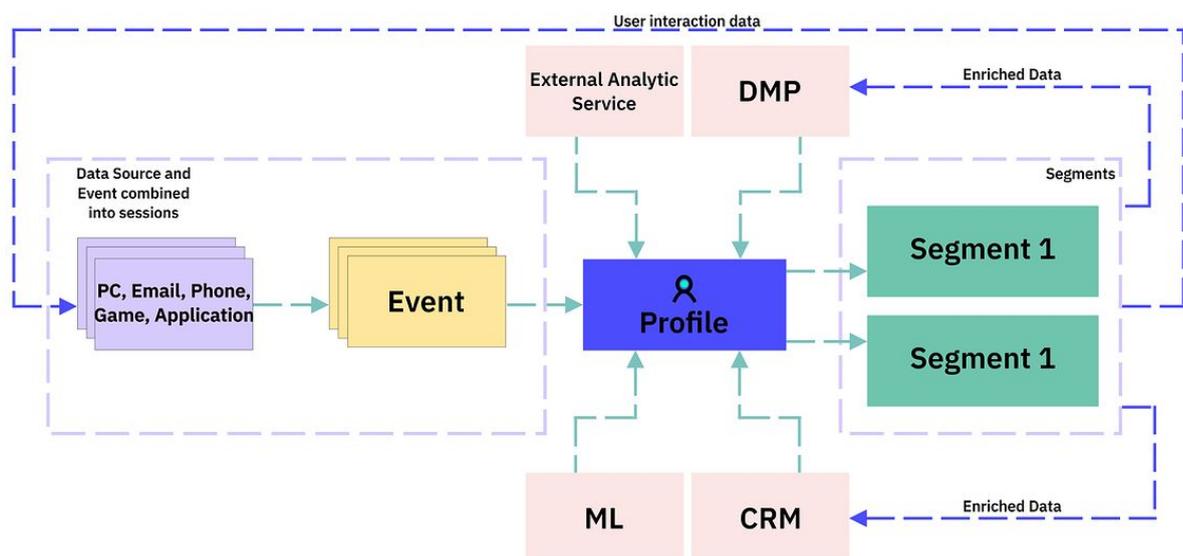
Starting work

The Uttu services collect and process information on user activity on web pages, in applications and in games. The processed data is used to personalize content, analyze user behavior, classify users into segments according to the metrics specified by the analyst, and build forecast models using this data. The data, reports and scenarios created in Uttu can be integrated and built into external CRM and advertising systems.

To get started, you need to integrate the SDK into the data source (on your website, into an application or game) and set up an Uttu project. If necessary, import existing data.

What happens to the data?

Uttu uses the concept of events to collect and process user information. It can accept events of arbitrary semantics and structure. Each event stores information about the time, type and additional data characterizing the parameters of the event. Regardless of the structure, events will be normalized, processed, saved and made available for building reports and segments.



Usually, an event represents information about the user's action or a change in some of their attribute. As Uttu processes events, it builds a presentation of each specific user and their behavior, collecting this data in a profile:

- Personal data (name, age, other personal data).
- Unique identifiers (e-mail address, cookies, fingerprint, platform and social networks identifiers, device identifiers).
- Purchases.
- Social profile (interests, reaction to content in social networks, friends graph).

An important feature of Uttu is its ability to work with large amounts of data in real time. The data is stored in column form, it can store trillions of events each with arbitrary number of parameters. Uttu can store parameters in the form of individual tuples or nested data structures for which you can build reports. See the [Data Storage](#) section for an explanation.

Most people use different devices, browsers and applications. It is not always easy to establish connection between a user profile and a specific individual. Uttu SDK uses special server queries to get the client user ID. Subsequently, this identifier is transmitted to the Uttu server with each query. Sometimes one person's data is stored across a number of profiles — as Uttu acquires more information, these profiles are combined.

Uttu meets all GDPR requirements for the storage of personal data.

Session

A session is a time-limited interaction between a user and an application with an integrated Uttu SDK. A session allows to bundle a sequence of actions performed by a user in the course of its duration. In the context of web applications, sessions are usually associated with HTTP sessions. For mobile and desktop applications they imply periods of user activity within the application.

Unique identifiers, session identifiers and client identifiers are used for single user's end-to-end attribution across a variety of sites and applications. Uttu also uses unique identifiers to combine profiles.

Segments

Uttu allows to use individual events, combinations or sequences of events, enriched data and analytics of predictive and forecast services to create an arbitrary number of user segments.

Uttu is capable of building segments based on multidimensional clustering, such as [RFM/RFD/RFE](#), and of building Look-a-like segments based on the existing ones.

Segments allow to analyze the effectiveness of specific traffic groups or to separate the profiles by interaction type. It is possible to select embedded segments or create new ones using the visual construction set, and apply them to the current or previously collected data, comparing parameters in different subgroups.

Each event or a sequence of events alongside an arbitrary set of rules become separate microsegments, which can be further combined into more complex segments or attributes. User markup by segments functions in real time.

Using a visual editor, the analyst can prepare a named segment and microsegment group for use in marketing. Using these segments for RTB or for analyzing the quality of traffic

does not require the marketer to configure and maintain them. See the [Scenario Construction Set](#) section for an explanation.

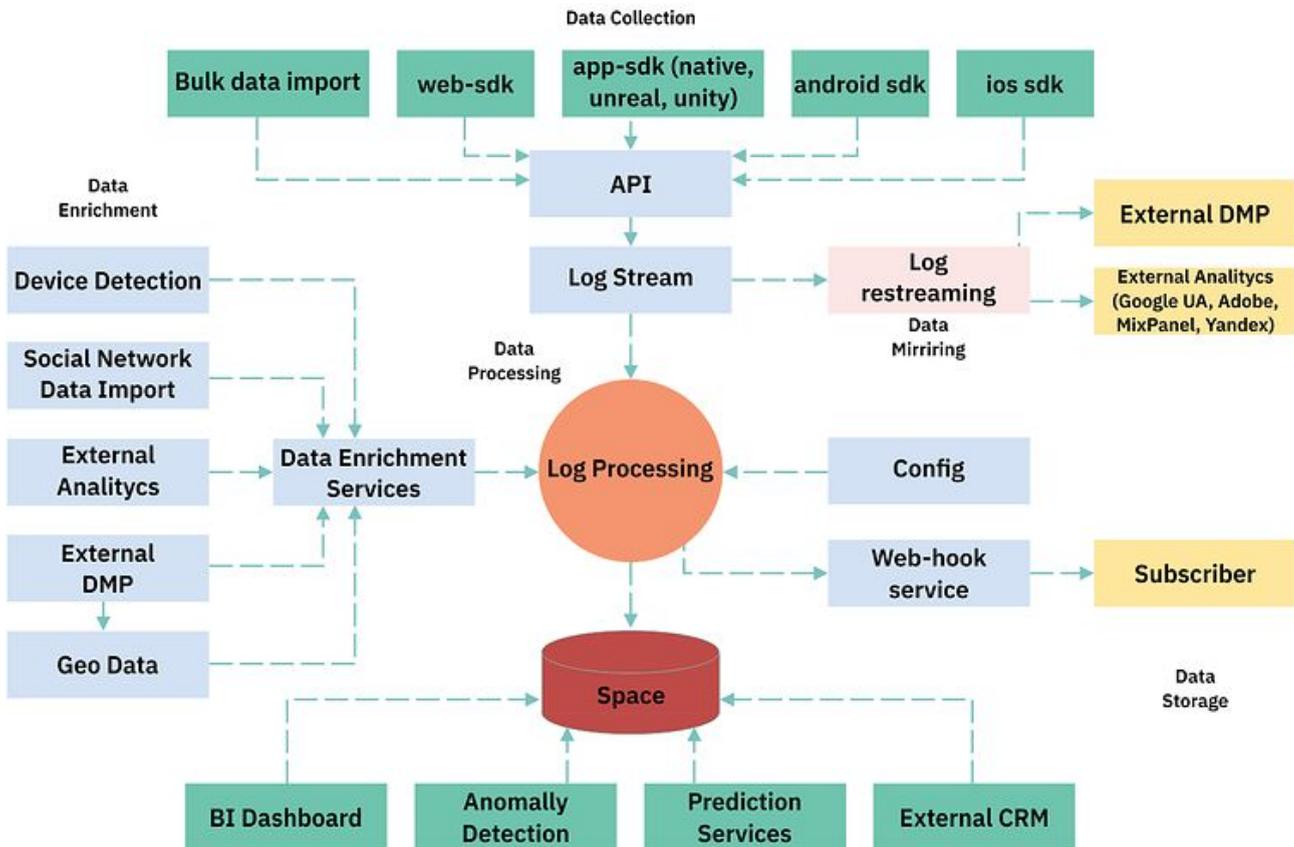
Data to build segments with can be sourced from the data collected by Uttu as well as from external sources, such as gaming and CRM databases, data from external analytics systems (Google Analytics, Adobe, MixPanel, etc.) or data loaded with the Bulk Import API.

Segments are used to analyze and select user groups, integrate with RTB, export data to external DMPs, or prepare renewable retargeting lists. In Uttu, this data is referred to as an “exported segment”, it can be formed in an impersonal form.

Structure

Working with Uttu implies the use of a web interface, client libraries and several APIs. The whole chain of data collection and processing can be divided into six main parts:

- **Data collection.** A combination of SDK and API for downloading data from various sources.
- **Mirroring.** A separate API class to redirect data to external systems and create services that respond to incoming raw data in real time.
- **Enrichment.** A combination of microservices and API extensions for searching, collecting, and structuring various additional information — for instance, the data pertaining to geography, devices, social profile, etc.
- **Processing.** The layer preparing the processed and enriched data for recording. It implements redirection of prepared data to external systems.
- **Storage.** Distributed secure infrastructure providing storage and direct access to data.
- **Analysis.** A set of data visualization tools and an infrastructure of forecast and analytical services providing predictive, statistical and behavioral analysis.



The data incoming to the public API goes to the streaming infrastructure. There, it is normalized and transmitted to a variety of services for data enrichment and forecast services. At this stage, segments are selected, scenarios are processed, and the data is written to the storage. The [“Architecture”](#) section provides further information on the technological stack.

SDK

Utzu includes a number of solutions for collecting data from various sources:

- Web SDK for tracking user activity on websites and in web applications.
- App SDK for tracking activity and collecting hardware information from desktop applications.
- Game SDK for integration with Unity, Unreal Engine.
- Mobile SDK for tracking user activity in mobile applications based on Android and IOS.

The App SDK allows to track the user’s behavior and his complex path from a webpage all the way to installation and launch of an application. By enabling the “Pass-Through Data” function, it is possible to transmit and receive additional data inside applications. This can be useful when a user registers on a website, then downloads an application or a game. Using “Pass-Through Data”, the user will be automatically logged into the application.

In order to integrate with e-commerce systems, you can use the Web SDK or install an extension. The systems currently supported include Shopify, Ecwid, Wechat, Wix, 3dcart, Wordpress, BigCommerce, Magento, Volusion, Big Cartel.

Bulk Import API

Bulk import is often used when migrating from a data collection and analytics system already being used to Uttu. It allows to automatically import data:

- directly from SQL (any ODBC-compatible source)
- NOSQL databases.
- using the REST API.

Uttu supports a number of modes of update from external sources:

- One-time download.
- Scheduled complete data reload.
- Incremental download.
- Data reaggregation.

The Bulk Import API, for example, is used to import data from the Protocol One [Hardware statistic](#) service. It permits to use the hardware information (operating system, GPU, CPU, display, the presence of camera, microphone, VR helmet) both for Uttu targeting and segmentation, and for the export of enriched data to external segments.

Data Hub

Data Hub is the [central hub](#) of both Qilin and the Protocol One DRM ecosystem, acting as an independent intermediary for the creation and storing of data common to all ecosystem participants.

P1 Data Hub ecosystem includes a variety of services and subsystems that send, enrich and aggregate data in one place. The Qilin Data Hub is mainly used to synchronize data on digital products for all ecosystem participants, and to aggregate and store sales data and royalty reports.

Uttu is integrated with Data Hub and is capable of sending and receiving data from the Qilin platform.

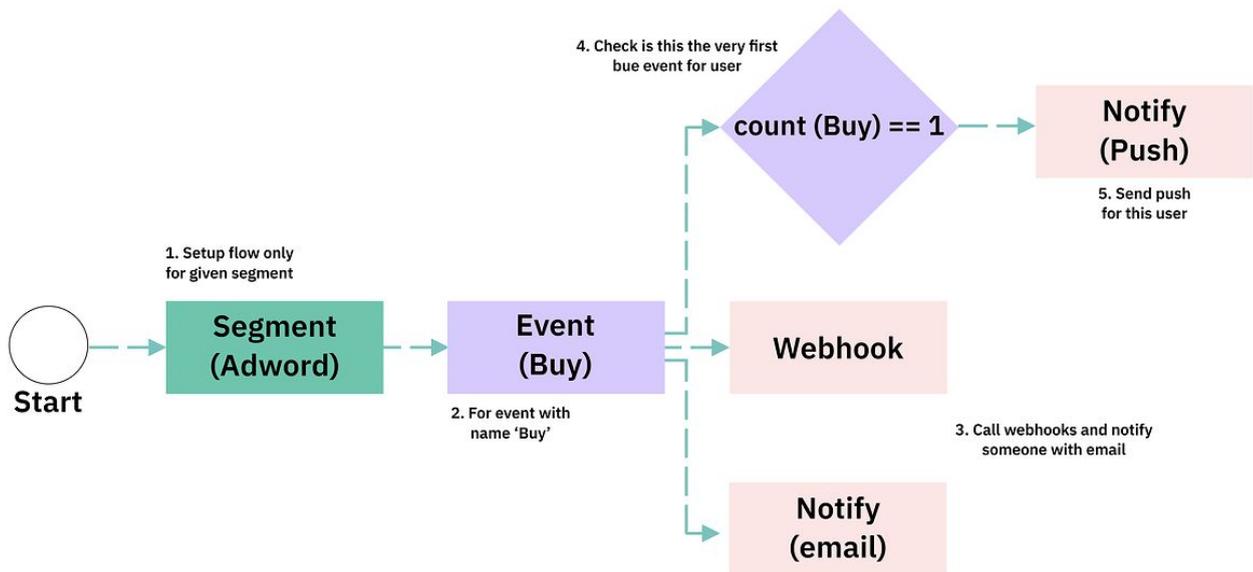
Analytics Tools

Automated traffic attribution and the grouping of it into segments in connection to marketing and advertising campaigns makes an important part of Uttu. It serves as a key channel of data on user behavior. Organic and search engine traffic, along with the data on users coming from social networks, are likewise collected into segments available by default. These segments serve as a base to create retargeting lists.

We plan to open a number of heuristic report services to analyze organics. At the moment Uttu lacks sufficient data to improve the accuracy of these reports.

Scenario Construction Set

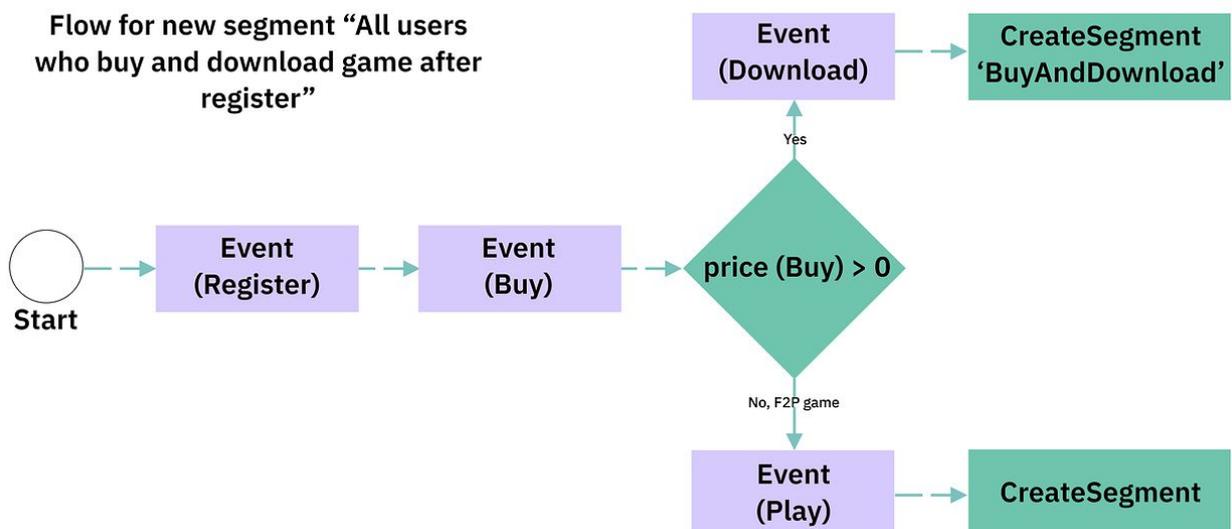
The Uttu Scenario Construction Set is used to create sequences of events, not unlike the “goals” in Google Analytics:



The diagram above describes the behavior of Uttu for users who have come through the Google Adwords advertising campaign. We send a notification, call the webhook in an external system and send e-mail — a receipt, for instance, — to every user who has finalized a purchased once.

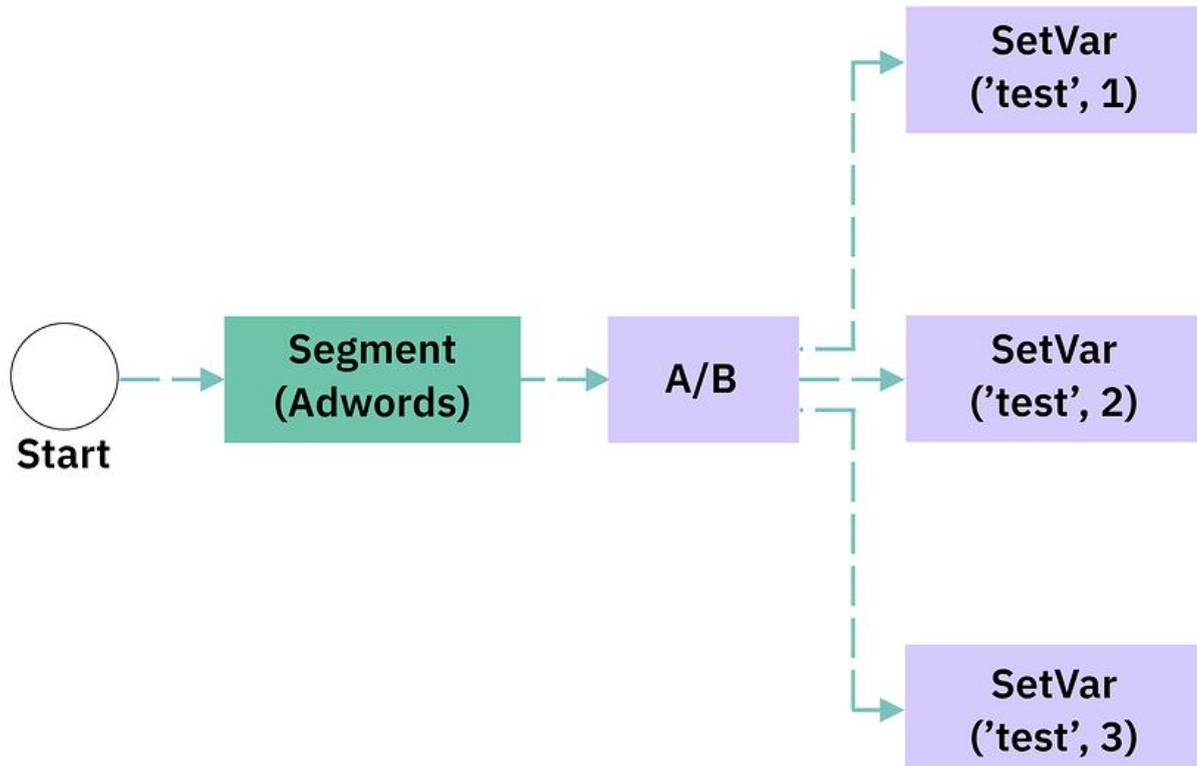
The Construction Set deals with two issues:

- the ability to program user interaction scenarios visually.
- creation of new segments based on a sequence of events.



Scenarios can be used to create new segments whenever necessary. These scenarios may include not just events but also conditions, other segments and scripts, time management statements.

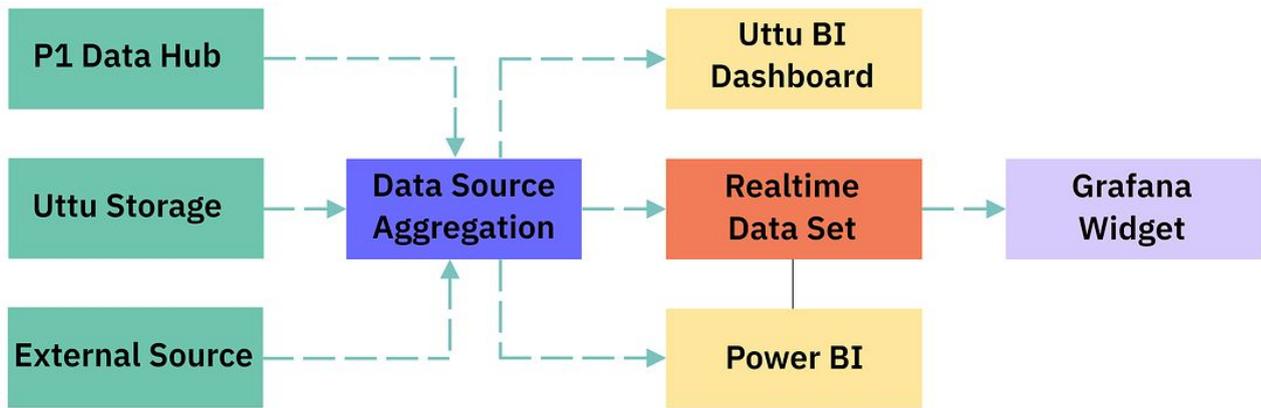
The Scenario Construction Set allows to create complex user interaction scenarios — for example, different reactions based on A/B tests:



Such scenario will be executed only for the specified segment. It sets different values for a variable that can be obtained from the target page. A/B tools can be used to send out a variety of notifications (e-mail, push notifications, SMS, etc.).

Toolbar

By default, Uttu supports a variety of toolbars, including Grafana and a solution of its own. You can integrate any data visualization and BI solutions that support data access using ODBC.



The sources of incoming data include Uttu data, common data from the P1 ecosystem and data from external sources integrated via connectors (databases, google docs, csv, data API, etc.). The aggregation interface allows a BI application to access the reports calculated in real time by flexibly customizing metrics and measurements built over the data combined from various sources.

Data visualization is handled by Uttu BI Dashboard. It supports:

- All widgets used for the analysis. Graphs and charts, tables, maps, filters and switches
- Widget connectivity. Applying a single filter to a group of widgets at once, drill-down, embedding widgets into another widget (for example, a table for each point of the graph).
- Flexible working with dataset. Mathematical functions, aggregations, sliding aggregations, window and text functions, complex calculations using Python and R.
- Variables: personal and user groups, session, custom
- Widget notifications, using e-mail, SMS and instant messengers. Based on threshold values, trends or patterns.
- Comments on specific widget elements and calendar events, using flexible visibility settings.
- Convenient sharing. Static pdf presentations for most widgets. It only takes three clicks to save and send it by e-mail.

Visual customization can be done using the construction set, or by the means of fragments from the template library.

The data selected in the widget by the customer can be immediately used to create a segment or cohort inside other elements of P1, including search for relevant audience to target advertising messages.

After the initial setup, the toolbar contains the following elements (see examples of specific reports in [Appendix 1](#)):

- End-to-end analytics for all channels and user interaction points
- Identification of traffic sources for advertising campaigns based on UTM tags, from social networks, search engines by keywords, organic traffic.

- Report templates for ad campaigns, traffic sources, devices, login pages, etc.
- Tools for segment construction, including look-a-like mechanics across the entire user base
- Ability to build arbitrary reports without having to plan data models in advance
- Reports on multichannel sequences

Multichannel sequences

In the simplest case, the analytics systems associate the target action performed by a user with one of these sources: the first, last, and last effective interaction.

The principle of multichannel sequences consists of tracking the entire interaction history (clicks and transitions) that resulted in the execution of the target action. The retrospective analysis window can be configured, by default it is 30 days long.

This approach allows to estimate each channel's contribution to the user performing the target action.

Scenarios for multichannel sequences help to find out:

- The role (weight) of a particular channel associated with the target action
- Means of conversion: main patterns that lead to the target action, time until the target action, the length of the sequence (the amount of interactions with the channels)

We are actively looking for a solution for effective interaction with e-commerce reports. This will enable us to compare expenses and revenues correctly and reliably, taking the entire chain of multichannel sequence into account.

How the different sources work:

- **First interaction.** Often used to track delayed conversions, when the first visit (the action) and the execution of the target action are separated by a significant amount of time. In this case, all visits from any sources will be associated with the original source of the first visit.
- **Last interaction.** This model does not account for the user's visit history. All actions are associated with the current transition source.
- **Last effective interaction.** Used by default, it allows to calculate conversions more accurately. All sources are divided into several lists, including significant sources (advertising, website links, search engines, social networks and mailing lists) and secondary sources (direct and internal transitions, transitions from saved pages). Visits from secondary sources are associated with the most significant primary source.

When creating reports, it is possible to select different interaction models.

Working with raw data

When working with complex scenarios or creating new forecast or analytical services, the analyst needs access to raw data. As a quick start, Uttu provides a download of customized sample of Jupiter's Notebook.

SQL support

Uttu supports an SQL-based declarative query language, matching the SQL standard in many cases. GROUP BY, ORDER BY, subqueries in the FROM, IN and JOIN sections, and scalar subqueries are all supported. Dependent subqueries and window functions are not supported.

Data is not only stored in columns, but also processed in vectors — that is, column fragments. This results in high efficiency of CPU usage.

The database supports arrays as well as nested structures and tuples. Events with an arbitrary set of parameters are stored in a natural way.

There is no database locking during data addition; it is possible to add and update data in real time.

Physical sorting of data by primary key allows to obtain data for a specific key value or a range of values with low latency, fewer than tens of milliseconds.

Due to low latency, it is not necessary to delay query execution or to prepare the response in advance. Queries are executed during loading of the user interface page, in online mode.

The SQL implementation supports approximate calculations:

1. Aggregate functions for approximate calculation of the number of different values, median and quantiles.
2. Support of queries based on partial (sampled) data, and approximate results.
3. The ability to perform aggregation for a limited amount of the first keys found. Following certain conditions for the distribution of keys across data, this function returns fairly accurate results using fewer resources.

For additional information, [visit the link](#).

Predictive services

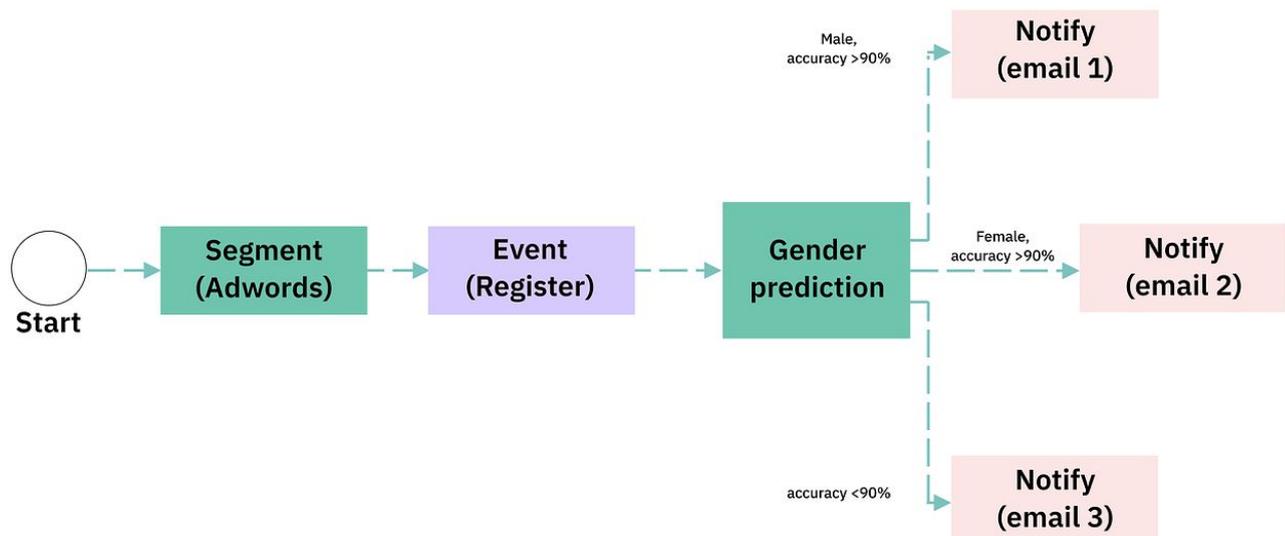
Uttu predictive services use proven mechanisms to solve highly applied tasks.

Prediction of individual user behavior. Suggestion system and prediction platform based on machine learning algorithms

- Prediction of LTV and lifetime in the game. For each player, it is possible to predict when he quits the game. For games using the free-to-play model, this allows to estimate the total revenue from a particular player. The first few days within the game are sufficient to make fairly accurate forecasts. Further on the prediction is refined with increased accuracy.
- “What if” scenarios. It is possible to predict the effectiveness of marketing communication or the outcome of changing the flow of such communications for each player or cohort.
- Individual suggestions of specific games, genres and settings, to create segments with high advertising efficiency.
- Individual pricing recommendations. By analyzing the players’ behavior patterns, it is possible to create individual discount offers, optimizing both the discount value and its precise timing.
- Optimal message delivery time. For each user, it is possible to predict the best time to deliver messages as far as conversion is concerned (considering the day of the week, time of day, time zone, individual user patterns)

Predictive service relies on a large array of data. This includes both the data belonging to a specific customer of the system, and the impersonal statistics of all users of the P1 ecosystem, along with enriched data coming from external sources. The format for data exchange within the system is based on Open GDE.

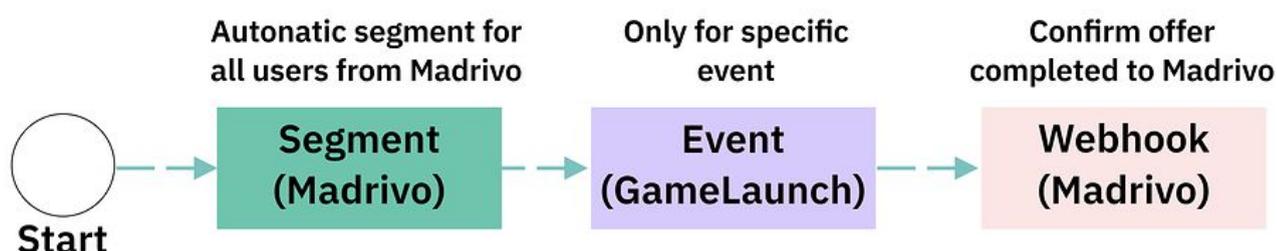
By default, predictive services do not create new events; they are only responsible for enriching the data and for building reports associated with them. Nonetheless, the processing infrastructure can use data from these services in form of conditions for the scenario construction set:



In the example above, all users registered by the means of the specified advertising campaign will receive customized e-mail notifications depending on their predicted gender, if the prediction accuracy exceeds 90%.

Ad networks

The Uttu Scenario Construction Set allows to automate S2S integration with a large amount of advertising platforms operating based on the CPA (Cost Per Action) and CPL (Cost Per Lead) models.



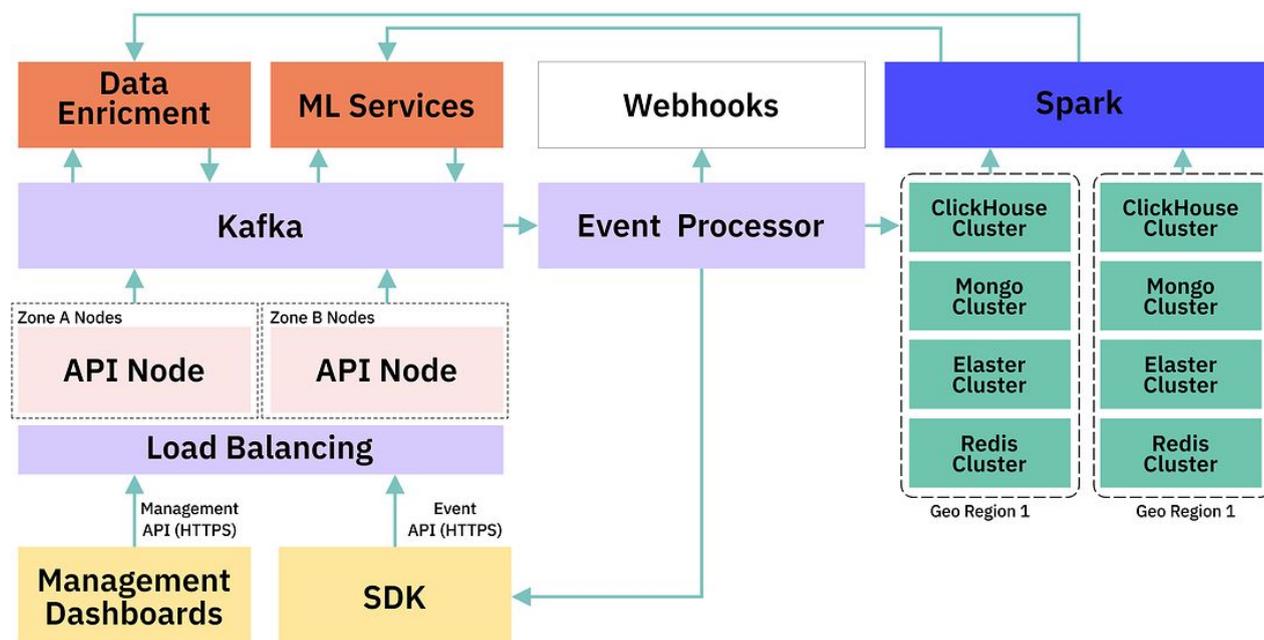
Support for Clickbooth, MaxBounty, admitad, alimama, W4, PeerFly, CPA Lead, Madrivo, qiyou, ClickDealer, everad, MonsterAds ad networks is available on Uttu out of the box.

The Uttu ad module allows the marketer to customize S2S integration with any ad network and start defining segments and reports at once without the help of engineers, using the visual construction set alone.

We are working to make the integration of Uttu with programmatic advertising networks based on Open RTB and Header Bidding as simple as possible. At the moment, Uttu can be integrated with RTBkit (Bidder-as-a-Service), Vanilla-rtb, platform.io and Adkernel as a DSP / DMP provider.

Our end goal is to implement an integrated Open RTB module and create infrastructure for rapid integration with SSP and AdExchange platforms and companies.

Architecture



Technology stack

- DNS routing based on geographic location.
- Nginx-based HTTPS
- Data enrichment API servers based on auto-scalable group in kubernetes
- Go and JavaScript ES6 as main languages.

Event handling

- Incoming events are written in Kafka Streams
- User data and segments are stored in Mongo and ClickHouse.
- Data processors (enrichment, ML, etc.) based on auto-scalable group in kubernetes read data from Kafka and perform segmentation, trigger calls, and handler calls in real-time.
- Events and segments are stored in Mongo (users, segments and scenarios), Clickhouse (raw data, reports, trends), Elasticsearch (search).

Suggested hardware

Choosing optimal hardware is no easy task, it depends on a large number of factors: the volume of incoming and stored data, the amount and complexity of integrations and the SLA requirements.

We give the following recommendations for unaided installation of Uttu on your own hardware, provided that the volume of your data is fairly small — up to 200 GB of compressed data per month (up to 300 million events per day):

- The fault tolerance policy suggest you to run at least 3 servers for each cluster group (Kubernetes, Clickhouse, Mongo, Redis, ElasticSearch, Kafka).
- Recommended server configuration:
 - Clickhouse:
 - 2 x Intel Xeon E5-2630 v4
 - 128 GB RAM
 - RAID-10 (or 5, 6, 50) on 8 or more SAS disks with rotation speed of 10k or above, or SSD
 - 10 Gbps network interfaces
 - ClickHouse recommendations:
<https://clickhouse.yandex/docs/en/operations/tips/>
 - Redis, Mongo, Kafka, ElasticSearch:
 - 2 x Intel Xeon E5-2630 v4
 - 256 GB RAM
 - RAID-10 on 4x400 GB SSD
 - Kubernetes:
 - Virtual machines with 8 GB RAM and 8 VCPU

Scalability

All components of the basic architecture are easily scaled horizontally in terms of the resources used for storing and processing data. Uttu stores large volume of data, and efficient and cost-effective handling of unpredictable workload is stands high among our priorities.

Data storage

[Yandex Clickhouse](#) serves as the main Uttu database. It is an open and very fast columnar database management system (DBMS) for online analytical processing (OLAP).

Unlike many data management methods, where a huge volume of raw data stored in a proprietary format is available as a “data lake” for any kind of query, ClickHouse for most cases offers instant results: the data is processed faster than it takes to create a query. It can process up to two terabytes of data per second, and is comparable with (and sometimes faster than) Vertica according to [the latest measurements](#). You can also browse these independent tests: [Spark](#), [MySQL](#) and [Redshift](#).

Data storage on disk

Many columnar DBMS (SAP HANA, Google PowerDrill) can only work inside RAM. This approach encourages greater expenses on hardware than is actually required for real-time analysis. ClickHouse is designed to work on ordinary hard drives, providing low storage cost

per gigabyte of data. However, SSD and additional RAM are also used to the full, if available.

Distributed query processing

Large queries are naturally parallelized for execution on all available processor cores, using all the necessary resources available on the server.

Nearly all columnar DBMS listed above do not support distributed query processing. In ClickHouse, data can be located on different shards. Each shard can be a group of replicas used to ensure fault tolerance. A query is executed on all shards in parallel. This is performed in transparency for the user.

Data protection

Processing and storing data with Uttu meets all requirements of [GDPR](#).

- Protocol One controls physical location of Uttu user data repositories.
- All external services and components that interact with Uttu user data meet all requirements of GDPR. A list of these services is available for viewing on a public page.
- Uttu personal data collection services always ask for user consent. Consent information is stored and the fact of consent can always be proved.
- Uttu does not collect personal data until consent is granted by the end user or the legal representative of the user who has not reached the legal age (16 years for Europe).
- End users are explicitly given the opportunity to acquaint themselves with the Uttu privacy policy. The policy is written in simple and understandable language.
- The Uttu website provides forms to request data deletion.
- Protocol One employees are properly trained to handle data in accordance with the GDPR
- All inconsistencies (or suspicions about their existence) can be reported to the contact e-mail address, or via the submission form available on the website.
- Any questions from end users concerning user data will be answered within 30 days from the date of request.
- For EU citizens, Uttu does not collect data on race, political opinion, religion, health, and sexual orientation, or data from where this information can be obtained indirectly.
- Protocol One has a public page that explains all aspects of compliance with the GDPR

Glutton Uttu

Uttu is a free and open platform —however, Protocol One uses an Uttu installation of its own to provide SAAS-based services and to integrate Uttu with other Protocol One services.

Most of the features are available in the open version of the platform, but some are only implemented and available in Glutton Uttu. These include the following:

Feature	Glutton	Uttu
Infrastructure with guaranteed SLA and support	Yes	No
Infrastructure meeting GDPR requirements	Yes	No*
External CRM integration	Yes**	
Predictor of user's social profile	Yes	No
Individual suggestions of specific games, genres and settings.	Yes***	No
Individual pricing recommendations.	Yes***	No
Predictor of optimal message delivery time.	Yes	No
RFM Analysis Tools	Yes	Partially
Integration with Protocol One Data Hub (CDP)	Yes	No
Integration of external segments from DMP	Yes	Partially****
Ready-to-work CPA integrations	Yes	No
Open RTB as a service	Yes	No

* Unaided installation of Uttu relies on the owner to meet the requirements of the laws concerning storage of personal data on his own. As for us, we provide the best practices and any help necessary to create such an infrastructure.

** Not all forms of integration developed by the Protocol One team are publicly available as a part of Uttu. Some of these are distributed on commercial terms.

*** Efficient operation of this service requires integration with Protocol One Data Hub.

**** Uttu features all APIs and tools necessary for such integration. Glutton Uttu offers a list of DMP providers out of the box, so that their data can be used immediately after registration.

Using unaided installation of Uttu is and will remain free. Listed below are examples of upcoming Glutton Uttu SAAS fees as a part of the Protocol One platform (the fees will be described in detail in WP 1.1 in March 2019):

- Free as long as the volume of stored data does not exceed 10 gigabytes, then \$10 for every 10 gigabytes of data per month.
- Using CPA / RTB integration: 0.1–2% of the revenue coming from the use of these tools.
- Separate pricing on prediction services.
- Paid installation and consulting for unaided installations of Uttu.

Disclaimer

This document does not represent a solicitation of investment or any other form of material support for the Protocol One project, Uttu or the Open Game Data Exchange. The purpose of this document is to provide a detailed and comprehensive description of Uttu.

Statements and other information of a declarative nature contained within this document must not be construed as direct assertions or promises unless they are expressly specified as such.

In the current edition, we actively collect feedback from analysts, marketers, video game developers and publishers, stores and platforms to improve the specifications and make the product we create easier and more convenient for all participants.

Appendix 1. Default settings

Right after installing Uttu, the data processing specialist is provided with a set of pre-installed and customized segments, events, and data.

Parameters

- **Sources**
 - Parameters based on UTM tags (Campaign, Ad Content, Source, Keyword, Channel Type)
 - Social networks (facebook.com, youtube.com, twitter.com, reddit.com, vk.com)
 - Search engines, websites and web pages from where the transition has occurred.
- **Technologies**
 - Type of device (PC, mobile phone, tablet, TV)
 - Device brand and model (for mobile phone, tablet, TV).
 - Operating system family and version (Windows, Android, IOS, etc.).
 - Browser (family, versions, supported features)
 - Availability and version of Flash, Java.
 - Display settings.
- **Content**
 - Login page
 - Visited page
 - Exit page
- **Geography**
 - Country
 - Region (administrative unit)
 - City (settlement)

Indicators

1. Number of sessions
2. Number of visitors (visitor is a user of the site within one browser)
3. Number of users
4. Average number of unique pages per session
5. Number of page views
6. Number of unique page views
7. Average session duration
8. Rejection rate
9. Number of conversions
10. Conversion rate (number of sessions with unique conversion to total number of sessions)

11. Real-time data slices (active sessions, conversions, sources, campaigns, pages, events)

Reports

1. Marketing activity (“Campaign - Source - Keyword”), based on the values from UTM tags. Indicators used:
 - Number of sessions.
 - Number of visitors.
 - Number of users.
 - Average session duration.
 - Rejection rate.
 - Number of conversions.
 - Conversion rate (number of sessions with unique conversion to total number of sessions).
 - Ad contents.
2. Report on traffic sources (general), grouped by categories.
 - Advertising campaigns (all traffic marked by UTM tags).
 - Traffic from search engines (traffic from known search engines).
 - Traffic from social systems (traffic from known social systems).
 - Traffic from websites (traffic from any other kinds of sites).
 - The second level of indicators consists of traffic sources, including those based on utm source.
3. Device report.
4. Browser report.
5. Login page report.
6. Visited page, by pages and events.
 - Unique page views.
 - Page views.
7. Report on the effectiveness of advertising campaign, based on traffic sources grouped by categories.
 - Advertising cost (when connected to an external data source).
 - Gross revenue from advertising.
 - Income from advertising.