Amazon Sustainability

amazon

AWS Customer Water Withdrawals Methodology

July 2026

Version 1.0

Table of Contents

List of Figures

Overview

This document outlines the methodology that governs the water withdrawals data in AWS' sustainability products. AWS may reasonably update this document upon material changes to the model logic and calculations.

The document is structured as follows:

1. **Scope** – drivers and scope, along with a look into the underlying standards governing the water withdrawals allocation model.

2. **System Boundary** – insight into attributable vs non-attributable water withdrawals, inclusions and exclusions.

3. **Input Data** – sources of data used in the water withdrawals model.

4. **Allocation** – how AWS water withdrawals are allocated to customer accounts.

1. Scope

Through its sustainability products, such as the AWS Sustainability service, AWS publishes the estimated water withdrawals associated with customers' AWS service usage. Customers use a wide array of AWS services across multiple AWS Regions making it complex to track and apportion water withdrawals based on workloads. The water withdrawals model aims to address this complexity, using the best available data and methodology. This document outlines the logic and calculations for version 1.0 of the model.

The water withdrawals model quantifies customer-specific water withdrawals associated with the use of AWS cloud services. It covers the full range of cloud products, from Amazon EC2 virtual and physical servers, to Amazon Simple Storage Service (S3), Amazon Elastic Block Store (EBS), Lambda, SageMaker, and all remaining AWS services.

The results are provided as cubic meters (m^3) and broken down by AWS Region and service. Data is reported on an annual basis by July of the following year; for example, 2025 data is reported by July 2026, and data is available since 2023. Historical water withdrawals data will be recast when significant methodology updates are introduced.

1.1 Definitions

The following terms are used throughout this methodology document:

Term	Description
Activity data	Activity data are quantitative measures of a level of activity that results in water withdrawals. It is a measure that models what is taking place during a given period (e.g., gigabytes of storage used, dollars spent).
Allocation	The action of apportioning water withdrawals of a given activity to a service based on an established weighting factor: e.g., usage-based ('physical allocation') or equivalent revenue-based ('economic allocation').
Assured data	Site-level water withdrawals data that has been independently reviewed by an external auditor.
AWS Sustainability service	An AWS service that helps customers track, measure, and review the estimated environmental impact associated from their AWS usage.
Cluster-level water withdrawals	The water withdrawals in cubic meters (m^3) for a particular AWS cluster. An AWS cluster can be an AWS Region (e.g. Frankfurt, Ohio), or an AWS CloudFront Edge Cluster (e.g. CF in North America, CF in South America).
Customer accounts (CA)	Customer accounts used to access cloud services. These accounts can be associated to different customers: 1. AWS internal teams - accounts owned by AWS service teams such as EC2, Lambda, etc. 2. Amazon businesses - accounts owned by teams in other Amazon business units beyond AWS. 3. External - accounts used by external customers.
Customer-segment-level water withdrawal	The water withdrawal in cubic meters (m^3) for a particular group of customers that are using a specific service.

Economic allocation	Revenue-based allocation method used as a fallback when physical usage data is unavailable or cannot be normalized across services and accounts. Economic allocation uses equivalent revenue (eqRevenue) as the weighting factor to apportion water withdrawals proportionally based on standardized usage-based revenue, allowing for consistent comparison across customer segments, time periods, and regions.
Edge sites	Types of data centers situated close to end customers, to enable ultra-low latency. Points of Presence (PoP), Cloudfront, and Transit sites are considered Edge.
Equivalent revenue (eqRevenue)	The equivalent revenue is the weighting factor leveraged when applying 'economic allocation'. It is a standardized measure of usage-based revenue that excludes pricing variation factors such as discounts, and other adjustments. It reflects service usage for each customer account, allowing for a consistent comparison across customer segments, time periods and regions.
Foundational service (FS)	AWS services that have dedicated server racks in AWS data centers.
m^3	Cubic meters of water withdrawn. The standard unit of measurement for water volume in AWS Sustainability products.
Non-foundational service (nFS)	AWS services with no dedicated server racks, which rely on foundational services to provide functionality to end customers.
Physical allocation	Usage-based allocation method applied when usage metrics are measurable (e.g., GB-year of storage used by a customer account). This is the preferred allocation method.
PPD	Planned Power Draw. The expected average power draw during peak usage conditions for a given server rack. The primary purpose of PPD values is to drive planning decisions to reserve sufficient power for each server rack deployed in AWS's data centers.
Ramp-up period	The initial phase of a new region or server rack deployment, defined as infrastructure less than 12 months old, during which utilization rates may be lower than mature infrastructure averages.
Reclaimed water	Water that has been treated and purified for reuse in data center operations (for example treated wastewater), typically for cooling purposes. Reclaimed water reduces reliance on freshwater sources and contributes to water stewardship goals. Reclaimed water is included in total withdrawals figures, consistent with GRI 303:2018, and tracked as a distinct source category.
Service-level water withdrawal	The water withdrawal in cubic meters (m^3) for a particular AWS service. This can be a footprint for a foundational or a non-foundational AWS service in any given cluster.
Top-down model	A water allocation approach where the model takes the total water

	withdrawals (i.e. starting from the top) and performs transformations to break down this single value into several logical segments. The model first separates the water withdrawals by cluster, then by server racks in each cluster, then by cloud services, and ultimately by customer accounts.
Water withdrawals	The total volume of water taken from any source (e.g., surface water, groundwater, rainwater, third-party/municipal supply, and reclaimed water) for data center operations, regardless of whether it is eventually returned to the environment. This definition aligns with the CDP Technical Note on Water Accounting and GRI 303:2018.

2 System Boundary

The system boundary defines what activities and related water withdrawals are accounted for in model calculations. The model allocates total water withdrawals from AWS data center operations to individual customer accounts based on their proportional cloud service usage, to ensure customers receive transparent, attributable water withdrawals data. To identify which activities are accounted for in model calculations, we define attributable and non-attributable water withdrawals.

“Attributable water withdrawals” is water withdrawn from any source to directly support the provision of cloud services for AWS customers – for example, water used in data center cooling systems, domestic water uses within operational data centers, and water used to fill a fire flow tank. The model includes water withdrawals from AWS owned, leased, or controlled data center facilities. This includes certain edge sites (e.g., CloudFront, Transit sites, Points of Presence). Primary activity data sources include utility bills for owned and leased facilities, and third-party operator reports for controlled facilities.

Conversely, water withdrawals not directly connected to AWS cloud service delivery are classified as “non-attributable water withdrawals.” The model excludes AWS warehouses and manufacturing facilities, AWS office buildings, and Customer-hosted sites (Amazon Cloud Extension sites, Embedded Points of Presence, Outpost locations) as they support general business operations rather than direct cloud service delivery.

The system boundary for water withdrawals is aligned in principle with the methodology used for customer carbon emissions allocation. The same criteria for attributability and the same facility scope apply, with differences limited to the nature of the environmental metric (cubic meters of water withdrawn vs. metric tons of CO_2 equivalent).

3 Input Data

The model calculates water withdrawals per cluster, on an annual basis, and subsequently allocates such volumes to customers based on their use of cloud services. This section outlines the sources of data used to define total water withdrawals for each AWS cluster. The customer allocation methodology is detailed in Section 4.

AWS estimates and assures water withdrawals data for all clusters, including cooling water across all facilities and domestic water at leased/owned facilities only. The water withdrawals model used to allocate to customers also includes domestic water use at AWS controlled data center facilities. Although these data are not included in 2025 assurance, as they represented less than 3% of the total AWS water footprint, they will be included from 2026 onwards. Domestic water use at these facilities is estimated based on facility size and the corresponding number of employees expected to work at each location.

AWS collects water withdrawals data at the site level across its data center portfolio. Data is sourced primarily from utility bills and third-party operator reports, with facility meters to fill in gaps. Standardized estimation procedures are applied where these sources are unavailable.

An independent auditor verifies all site-level water withdrawals data annually. Site-level data is aggregated at a cluster level, which is what is used in the water withdrawals allocation model.

4 Allocation

The water withdrawals described in section 3 are processed in an allocation model that calculates customers' water withdrawals associated with the use of AWS's cloud services. AWS prioritizes 'physical allocation' – i.e., usage-based allocation – and considers 'economic allocation' – i.e., based on equivalent revenue – as a fall-back option. Meaning, if physical data is not reasonably available, we rely on economic factors.

The model takes water withdrawals associated with each AWS cluster and performs a series of transformations to break them down into several logical segments. Conceptually, the model works using the following logical transformation workflow:

1. Allocate cluster-level withdrawals to the server racks in the cluster using the server racks' power draw.

> For instance, suppose a cluster is allocated 500 m^3 and has a total power draw of 1000 kVA. If the total power draw of server rack 1 is 600 kVA, it will get allocated 60% (600/1000) of the water withdrawals or 300 m^3.

2. Allocate water withdrawals associated with server racks to AWS cloud services based on utilization of server racks, accounting for inter-dependencies between:
    a. 'Foundational services', AWS services that have dedicated racks in a data center.
    b. 'Non-foundational services' AWS services with no dedicated racks, which rely on other cloud services to provide functionality to end customers.

> Following our previous example, the model identifies that server rack 1 has a total usage of 1000 GB-year in that cluster. Scanning through the usage records the model observes that 250 GB-year is being used up by cloud service 1. The model assigns 25% (250/1000) of server rack 1's withdrawals to cloud service 1, which is 75 m^3 (25% * 300 m^3).

3. Allocate water withdrawals associated with each cloud service to individual customer accounts. For foundational services, this is done with physical allocation, while we use economic allocation for non-foundational services.

> Using the same example, the model identifies that cloud service 1 (which had a usage of 250 GB-year) is used across several customer accounts. Customer account 1 used 125 GB-year and is therefore assigned 50% (125/250) of withdrawals, which equals to 37.5 m^3 (50% * 75 m^3).

Figure 1 – Conceptual workflow for the allocation of water withdrawals to AWS customers

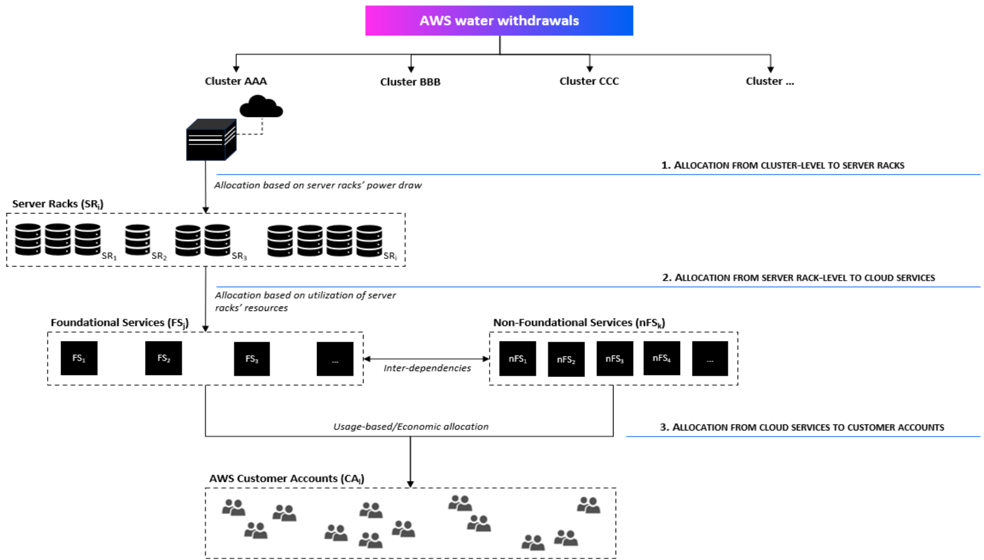

Temporal granularity of site-level water withdrawal data

The methodology prioritizes actual water withdrawal data where available. Most operators of AWS controlled facilities currently report withdrawal data on an annual basis. Because this data is annual, the allocation methodology currently does not capture seasonal variability in facility water use. The allocation remains strictly proportional to each customer's total resource usage, regardless of when within the year that usage occurs.

The methodology does not penalize customers based on the water characteristics of the facility where their workloads are placed within a cluster. Customers do not choose their facility placement. Allocation outcomes reflect only proportional resource usage, not site-specific characteristics beyond a customer's control.

As more granular reporting becomes available from operators of AWS controlled facilities, it will improve the temporal resolution of the allocation while preserving fair and consistent treatment across all customer accounts.

4.1 Cluster to Server Racks

Intent	Allocate cluster-level water withdrawals to the underlying server racks. Applicable for each AWS cluster.
Input	Cluster-level water withdrawals (m^3/year)
Output	Water withdrawals per server rack (m^3/year)
Type of allocation	Physical allocation based on PPD

Figure 2 – Allocation of cluster-level withdrawals to the underlying server racks[1]

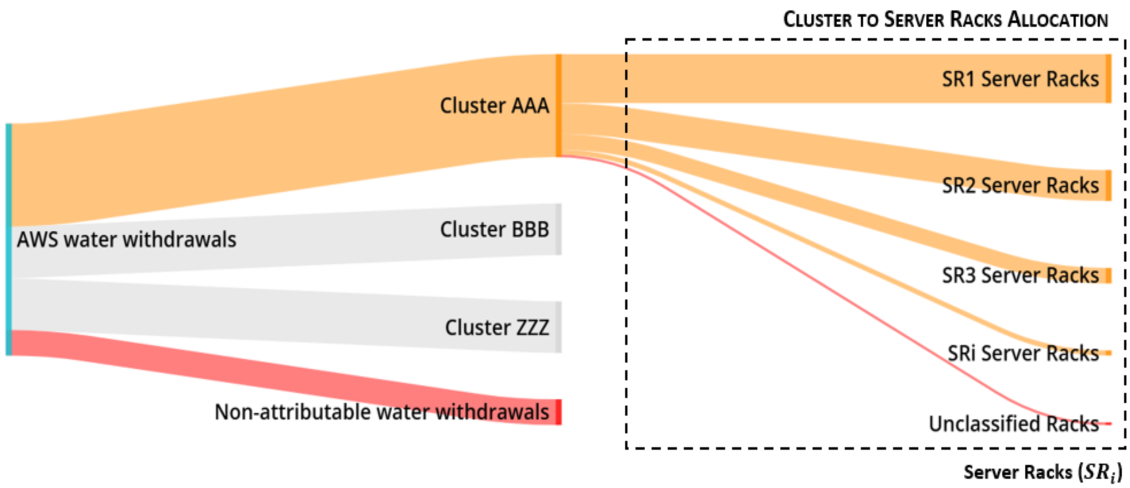

In this first allocation step, the model allocates cluster-level water withdrawals to the underlying server racks in the cluster. The weighting factor for the allocation is the share of the total server racks' PPD in the cluster. Such share is obtained as follows: for each server rack type (e.g., EC2.C5), the model takes the number of server racks ('population'), multiplies it by the average PPD of that specific server rack type, and divides it by the total server racks' PPD in the cluster. This produces the weighting factor, i.e., a percentage based on the server rack type distribution of that cluster. Let $\%\ SR_i$ be the share of power draw attributable to server racks type *i*, n_{SR_i} their population in the cluster, and PPD_{SR_i} the respective PPD of the server rack type *i*:

$$\%\ SR_i = \frac{n_{SR_i} * PPD_{SR_i}}{\sum_i (n_{SR_i} * PPD_{SR_i})} \quad \text{Eq 1}$$

where:

$$i = type\ of\ server\ rack\ [EC2\ (EC2.C5, EC2.C6N, ..); S3(standard, glacier, ..); ...]$$

$$\sum_i (n_{SR_i} * PPD_{SR_i}) = (n_{SR_{EC2.C5}} * PPD_{SR_{EC2.C5}}) + (n_{SR_{EC2.C6N}} * PPD_{SR_{EC2.C6N}}) + (n_{SR_{S3\ standard}} * PPD_{SR_{S3\ standard}}) + \cdots = Total\ server\ racks'\ PPD\ in\ the\ cluster$$

The water withdrawals associated with a given server rack type *i* in the cluster is:

$$W_{SR_i} = (W_{cluster} * \%\ SR_i) \quad \text{Eq 2}$$

For instance, suppose a cluster has 500 m^3 of associated water withdrawals and has a total planned power draw of 1000 kVA. If there is a total of 100 racks in such cluster with 80 SR_1 consuming a total of 150 kVA, 15 SR_2 consuming 50 kVA, and 5 SR_3 consuming 800 kVA, the model determines that:

- 15% of the cluster resource allocation should be assigned to SR_1, which is 75 m^3
- 5% of the cluster resource allocation should be assigned to SR_2, which is 25 m^3
- 80% of the cluster resource allocation should be assigned to SR_3, which is 400 m^3

1. The Sankey diagram is for illustrative purposes only and does not reflect actual split of water withdrawals across clusters, server racks, services and customers.

Note: usage and withdrawals in this example are not based on real data

4.2 Server Rack to Cloud Services

Intent	Allocate water withdrawals associated with server racks to the cloud services using them. Applicable for each AWS cluster.
Input	Water withdrawals per server rack (m^3/year)
Output	Water withdrawals per cloud service (m^3/year)
Type of allocation	Usage-based allocation

Figure 3 – Allocation of withdrawals to cloud services[2]

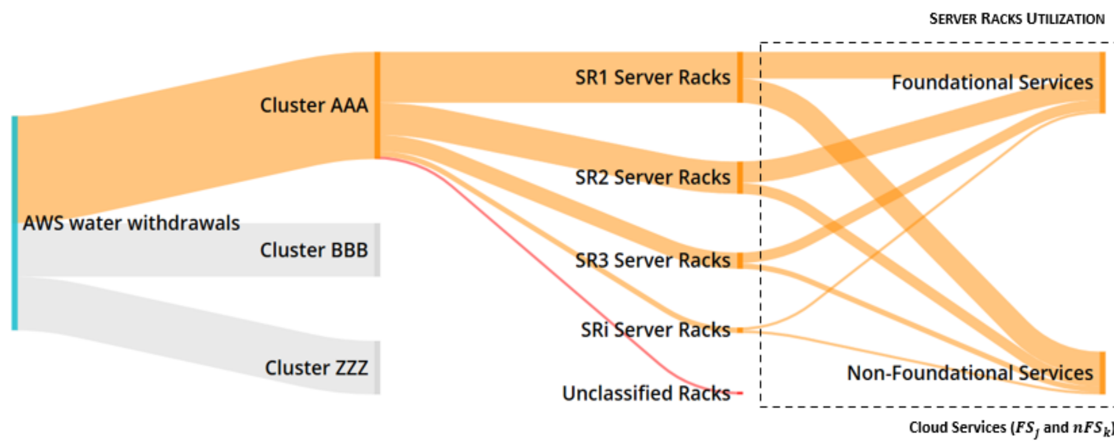

In this allocation step, the model defines a unit of usage for each server rack that is used to calculate the % of usage from a given cloud service – foundational services *j* (*% FS_j* Use of *SR_i*) and non-foundational services *k* (*% nFS_k* Use of *SR_i*) – out of the total usage of a server rack *i*.

The units of usage for each server rack are determined and maintained by each AWS team managing the respective foundational service: e.g., normalized Instance-Hour (NIH) for EC2 rack types, gigabytes (GB) of storage for S3 rack types, etc.

Note: the withdrawals associated with a type of server rack (W_{SR_i}) – e.g., EC2 types of server racks – is not equal to the withdrawals associated with the respective foundational service (W_{FS_j}) – e.g., EC2 – unless the server rack supports 'bare metal' types of instances, i.e., physical servers provided to customers rather than in a virtualized fashion. Outside of bare metal instances, server racks also support non-foundational services; thus, a portion of the water withdrawals associated with such racks is allocated to non-foundational services.

Essentially:

$$\Sigma_i W_{SR_i} \neq \Sigma_j W_{FS_j}$$

2. The Sankey diagram is for illustrative purposes only and does not reflect actual split of water withdrawals across clusters, server racks, services and customers.

The share % of usage of a server rack resources from a given cloud service is calculated as follows:

For **foundational services** (*FS*$_j$)

$$\% \, FS_j \; Use \; of \; SR_i = \frac{FS_j \; Use \; of \; SR_i}{Tot \; Use \; of \; SR_i} \qquad \text{Eq 3}$$

For **non-foundational services** (*nFS*$_k$)

$$\% \, nFS_k \; Use \; of \; SR_i = \frac{nFS_k \; Use \; of \; SR_i}{Tot \; Use \; of \; SR_i} \qquad \text{Eq 4}$$

where:

$$i \; = \; type \; of \; server \; racks \; [EC2; \; S3; \; \ldots]$$
$$j \; = \; type \; of \; foundational \; service$$
$$k \; = \; type \; of \; non \; foundational \; service$$
$$Tot \; Use \; of \; SR_i = \sum\nolimits_j FS_j \; Use \; of \; SR_i + \sum\nolimits_k nFS_k \; Use \; of \; SR_i$$
$$\sum\nolimits_j FS_j \; Use \; of \; SR_i \; = FS_1 \; Use \; of \; SR_i + FS_2 \; Use \; of \; SR_i + FS_3 \; Use \; of \; SR_i + \cdots$$
$$\sum\nolimits_k nFS_k \; Use \; of \; SR_i \; = nFS_1 \; Use \; of \; SR_i + nFS_2 \; Use \; of \; SR_i + nFS_3 \; Use \; of \; SR_i + \cdots$$

There are aggregations occurring in the model, which adjust the server racks usage (*% FS*$_j$ Use of *SR*$_i$ and *% nFS*$_j$ Use of *SR*$_i$) to account for different usage pathways/scenarios of foundational and non-foundational services.

Finally, water withdrawals associated with foundational services *j* and non-foundational service *k* are obtained as:

For **foundational services** (*FS*$_j$)

$$W_{FS_j} = \sum_i \left[\left(\% \, FS_j \; Use \; of \; SR_i\right) * \left(W_{SR_i}\right) \right] \qquad \text{Eq 5}$$

$$= \left(\% \, FS_j \; Use \; of \; SR_{EC2}\right) * \left(W_{SR_{EC2}}\right) + \left(\% \, FS_j \; Use \; of \; SR_{S3}\right) * \left(W_{SR_{S3}}\right) + \cdots$$

For **non-foundational services** (*nFS*$_k$)

$$W_{nFS_k} = \sum_i \left[\left(\% \, nFS_k \; Use \; of \; SR_i\right) * \left(W_{SR_i}\right) \right] \qquad \text{Eq 6}$$

$$= \left(\% \, nFS_k \; Use \; of \; SR_{EC2}\right) * \left(W_{SR_{EC2}}\right) + \left(\% \, nFS_k \; Use \; of \; SR_{S3}\right) * \left(W_{SR_{S3}}\right) + \cdots$$

where:

$$i \; = \; type \; of \; server \; racks \; [EC2; \; S3; \; ..]$$
$$j \; = \; type \; of \; foundational \; service$$
$$k \; = \; type \; of \; non - foundational \; service$$

$$W_{SR_i} = \; (W_{cluster} * \%SR_i);$$
$$i.e., water \; withdrawals \; associated \; with \; server \; racks \; type \; i \; in \; the \; cluster$$

Following our previous example, suppose the model assigned 400 m^3 to server racks 3 in a cluster. The model identifies that server racks 3 have a total usage of 1000 GB-year in that cluster. Scanning through the usage records the model observes that:

- 50 GB-year is being used up by cloud service 1
- 700 GB-year is being used up by cloud service 2
- 250 GB-year is being used up by cloud service 3

This means that the service utilization is:

- 5% to cloud service 1, equal to 20 m^3
- 70% to cloud service 2, equal to 280 m^3
- 25% to cloud service 3, equal to 100 m^3

Note: usage and withdrawals in this example are not based on real data.

Utilization cases

The allocation logic accounts for varying levels of resource utilization across AWS infrastructure, depending on the maturity of the server racks, and the cluster. The logic enables more equitable allocation of withdrawals in those cases where the service infrastructure is in its early phases ('ramp-up' period). For infrastructure in the ramp-up period (regions and racks <12 months old), the model compares utilization rates against mature infrastructure (regions and racks ≥12 months old) of the same rack type to identify any utilization anomalies. When utilization is lower than the mature infrastructure average, the resulting additional withdrawals are distributed as a global overhead across all services using that specific resource rack type. This approach prevents overburdening of services and their users in a particular cluster by sharing the withdrawals associated with the underlying operational overhead in the ramp-up of new clusters/server racks pro rata across services and their users at global level.

4.3 Cloud Services to Customer Accounts

Intent	Allocate water withdrawals associated with cloud services to individual customer accounts. Applicable for each AWS cluster.
Input	Water withdrawals per cloud service (m^3/year)
Output	Water withdrawals per customer account (m^3/year)
Type of allocation	Usage-based allocation for foundational services; economic (revenue-based) allocation for non-foundational services.

Figure 4 – Allocation of withdrawals to customer accounts[3]

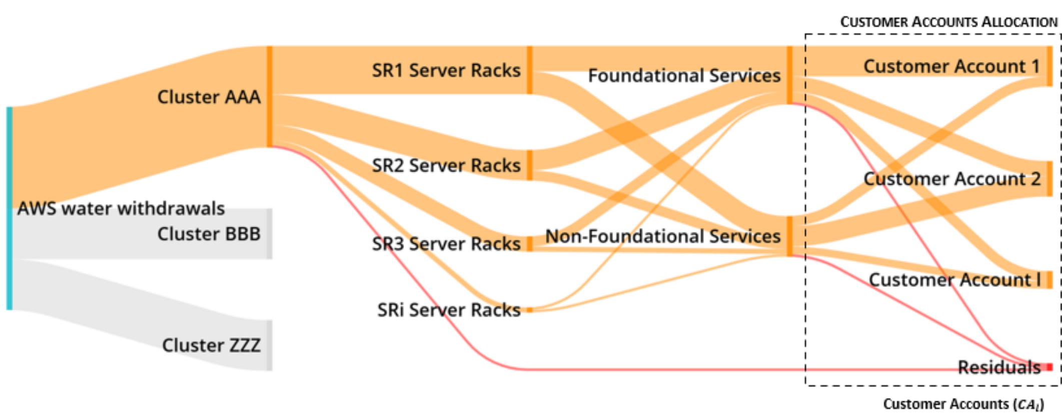

3. The Sankey diagram is for illustrative purposes only and does not reflect actual split of water withdrawals across clusters, server racks, services and customers.

The type of allocation method used to allocate cloud services' water withdrawals to customer accounts depends on the data available to track the usage of such services from individual accounts. We prioritize 'physical allocation' when usage metrics are available (e.g., customer account l using X GB in the accounting year using S3 services in a given cluster). We consider 'economic allocation' as fall-back option when it is not possible to determine a normalized weighting factor of usage across services and customer accounts, or usage data is not complete.

4.3.1 Foundational Services

The share of withdrawals attributed to foundational services in a given cluster is split among the customer accounts hosting workloads on that cluster, via usage-based allocation. Essentially, each customer account *l* receives a share of foundational services withdrawals proportional to its use of a given foundational service *j* out of total usage of such foundational service for that cluster, using the units of usage identified in section 4.2.

Let *% CA_l* Use of *FS_j* be the share of usage of foundational service *j* from customer account *l*:

$$\% \, CA_l \; Use \; of \; FS_j = \frac{CA_l \; Use \; of \; FS_j}{Tot \; Use \; of \; FS_j} \qquad \text{Eq 7}$$

Where:

$$j = type \; of \; foundational \; service$$
$$l = customer \; account \; ID \; [XXXXXXXXXXXX; YYYYYYYYYYYY; \ldots]$$
$$CA_l \; Use \; of \; FS_j = Use \; of \; foundational \; service \; j \; by \; customer \; account \; l \; in \; a \; given \; cluster$$
$$Tot \; Use \; of \; FS_j = \sum\nolimits_l CA_l \; Use \; of \; FS_j$$

Therefore, the customer account *l* withdrawals associated with the use of a foundational service *j*, ($W_{CA_l}^{FS_j}$) are:

$$W_{CA_l}^{FS_j} = (\% \, CA_l \; Use \; of \; FS_j) * \left(W_{FS_j}\right) \qquad \text{Eq 8}$$

Where W_{FS_j} are the water withdrawals associated with the foundational service *j* in the cluster, as calculated in section 4.2.

Let's assume that in a cluster, cloud service 1 has associated withdrawals equal to 300 m^3 after step 4.2, and that the service is used across three customer accounts:

Customer account 1, 400 GB-year

Customer account 2, 200 GB-year

Customer account 3, 200 GB-year

The total usage of cloud service 1 is 800 GB-year. Therefore, the model allocates withdrawals to these customer accounts based on their share of usage of cloud service 1:

Customer account 1 receives 50% of cloud service 1 withdrawals, i.e. 150 m^3

Customer account 2 and 3, receive 25% of cloud service 1 withdrawals each, i.e. 75 m^3

Note: usage and withdrawals in this example are not based on real data.

4.3.2 Non-Foundational Services

The share of a cluster's withdrawals attributed to non-foundational services is allocated to customer accounts via economic allocation. An account-level equivalent revenue (*eqRevenue*) is the established weighting factor for the allocation. The equivalent revenue provides a pure reflection of service usage for each customer account, allowing for a consistent comparison across customer segments, time periods and regions[4]. Each customer account *l* receives a share of withdrawals proportional to its associated equivalent revenue for a given non-foundational service *k*. Let $\%\ CA_l$ Use of nFS_k be the share of usage of non-foundational service *k* from customer account *l*:

$$\%\ CA_l\ Use\ of\ nFS_k = \frac{CA_l\ eqRevenue\ for\ nFS_k}{Tot\ eqRevenue\ for\ nFS_k} \qquad \text{Eq 9}$$

Where:

$$k = type\ of\ non\ foundational\ service$$
$$l = customer\ account\ ID\ [XXXXXXXXXXXX; YYYYYYYYYYYY; \ldots]$$
$$Tot\ eqRevenue\ for\ nFS_k = \sum_l CA_l\ eqRevenue\ for\ nFS_k$$

Therefore, the customer account *l* withdrawals associated with the use of a non-foundational service *k*, $W_{CA_l}^{nFS_k}$ are:

$$W_{CA_l}^{nFS_k} = (\%\ CA_l\ Use\ of\ nFS_k) * \left(W_{nFS_k}\right) \qquad \text{Eq 10}$$

Where W_{nFS_k} are the water withdrawals associated with the non-foundational service *k* in the cluster, as calculated in section 4.2.

Let's assume that in a cluster, the non-foundational service 1 (nFS_1) has associated withdrawals equal to 200 m^3 after step 4.2. The model scans the revenue stream associated with the use of nFS_1 and finds that, out of a total of $100k equivalent revenue:

- Customer account 1, $60k, 60% of total equivalent revenue
- Customer account 2, $30k, 30% of total equivalent revenue
- Customer account 3, $10k, 10% of total equivalent revenue

The model allocates the withdrawals associated with the cloud service nFS_1 (200 m^3) based on the customer accounts' % of equivalent revenue:

- Customer account 1, 120 m^3
- Customer account 2, 60 m^3
- Customer account 3, 20 m^3

Note: equivalent revenue and withdrawals in this example are not based on real data.

4. The equivalent revenue is a standardized measure of usage-based revenue that excludes pricing variation factors such as discounts, and other adjustments.

4.3.3 Aggregating All Together

The water withdrawals associated with a customer account *l* from the use of cloud services in a given cluster for the accounting year are:

$$W_{CA_l} = \sum_j W_{CA_l}^{FS_j} + \sum_k W_{CA_l}^{nFS_k} \qquad \text{Eq 11}$$

Where:

$$j = type\ of\ foundational\ service$$
$$k = type\ of\ non\ foundational\ service$$
$$l = customer\ account\ ID\ [XXXXXXXXXXXX;\ YYYYYYYYYYYY;\ \ldots]$$
$$\sum_j W_{CA_l}^{FS_j} = CA_l\ total\ water\ withdrawals\ from\ use\ of\ foundational\ services$$
$$\sum_k W_{CA_l}^{nFS_k} = CA_l\ total\ water\ withdrawals\ from\ use\ of\ non-foundational\ services$$

4.3.4 Allocation of Residual Water Withdrawals

The overarching principle in the water withdrawals customer allocation model is the 'preservation of water withdrawals', i.e., allocating the entirety of water withdrawals in the system boundary across clusters to customer accounts. For every accounting year:

$$\sum_m W_{Cluster_m} = \sum_{l,m} W_{CA_l}^{cluster_m} \qquad \text{Eq 12}$$

Where:

$$W_{CA_l}^{cluster_m} = CA_l\ total\ withdrawals\ associated\ with\ cluster\ m$$

As the final step, the model checks for any residual water withdrawals, and allocates them to customer accounts. Residual water withdrawals can stem from, for example, incomplete mapping of server rack types or services to accounts, or model processing errors. Following the preservation principle, the model reallocates the residual water withdrawals proportionally – based on customer accounts' shares of withdrawals – using two different level of aggregation:

- **Cluster-level**, if the root cause for the residual withdrawals can be traced back to a specific cluster. In this case, the residual withdrawals are reallocated pro-rata among customer accounts associated with that specific cluster within each scope.

- **Global-level**, if the root cause for the residual withdrawals is not directly attributable to a cluster or is not clearly identifiable. In this case, the residual withdrawals are reallocated pro-rata among all customer accounts at global level within each scope.

AWS conducts regular quality assessments to verify the residual allocation method is working as expected.

5 Bibliography

Standards and Frameworks

- CDP Technical Note on Water Accounting. CDP. https://cdn.cdp.net/cdp-production/cms/guidance_docs/pdfs/000/001/561/original/CDP-technical-note-water-accounting-definitions.pdf

- GRI 303: Water and Effluents 2018. Global Reporting Initiative. https://www.globalreporting.org/standards/standards-development/topic-standard-project-for-water/

Related AWS Documentation

- Customer Carbon Emissions Methodology Document, Version 3.0, Amazon Web Services https://sustainability.aboutamazon.com/aws-customer-carbon-footprint-methodology.pdf