



FOCUS

FinOps Open Cost and Usage Specification

Version

Publication version 1.2

Copyright © 2023-2025 - FinOps Open Cost and Usage Specification (FOCUS) a Series of the Joint Development Foundation Projects, LLC. Joint Development Foundation [trademark](#), and document use rules apply.

Status of This Document

This section describes the status of this document at the time of its publication.

This is a published release of the FinOps Open Cost and Usage Specification.

This document was produced by a group operating under the Joint Development Foundation Projects agreement. FOCUS maintains a public list of any patent disclosures made in connection with the deliverables of the group; [that page](#) also includes instructions for disclosing a patent. An individual who has actual knowledge of a patent which the individual believes contains Essential Claim(s) must disclose the information.

Document Use License

Copyright (c) Joint Development Foundation Projects, LLC, FinOps Open Cost and Usage Specification (FOCUS) Series and its contributors. The materials in this repository are made available under the Creative Commons Attribution 4.0 International license (CC-BY-4.0), available at <https://creativecommons.org/licenses/by/4.0/legalcode>.

Shield: 

This work is made available under: [Creative Commons Attribution 4.0 International License](https://creativecommons.org/licenses/by/4.0/legalcode).



THESE MATERIALS ARE PROVIDED "AS IS." The parties expressly disclaim any warranties (express, implied, or otherwise), including implied warranties of merchantability, non-infringement, fitness for a particular purpose, or title, related to the materials. The entire risk as to implementing or otherwise using the materials is assumed by the implementer and user. IN NO EVENT WILL THE PARTIES BE LIABLE TO ANY OTHER PARTY FOR LOST PROFITS OR ANY FORM OF INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES OF ANY CHARACTER FROM ANY CAUSES OF ACTION OF ANY KIND WITH RESPECT TO THIS DELIVERABLE OR ITS GOVERNING AGREEMENT, WHETHER BASED ON BREACH OF CONTRACT, TORT (INCLUDING NEGLIGENCE), OR OTHERWISE, AND WHETHER OR NOT THE OTHER MEMBER HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This document is governed by the Patent Policy Option 4: W3C Mode: See [project charter](#).

Abstract

FOCUS is an open specification for billing data. It defines a common schema for billing data, aligns terminology with the FinOps Framework and defines a minimum set of requirements for billing data. The specification provides clear guideline for billing data generators to produce FinOps-serviceable data. The specification enables FinOps practitioners to perform common FinOps capabilities such as chargeback, cost allocation, budgeting and forecasting etc. using a generic set of instructions, regardless of the origin of the FOCUS compatible dataset.

Working Group

Maintainers

Thanks to the following FOCUS Maintainers for their leadership and contributions to the FOCUS Release **v1.2** specification.

- Alex Hullah (Goldman Sachs)
- Christopher Harris (Datadog)
- Irena Jurica (Neos)
- Joaquin Prado (FinOps Foundation)
- Karl Kraft (Walmart)
- Larry Advay (Twilio)
- Michael Flanakin (Microsoft)
- Mike Fuller (FinOps Foundation)
- Riley Jenkins (Domo)
- Shawn Alpay (FinOps Foundation)
- Udam Dewaraja (StitcherAI)

Contributors

Thanks to the following FOCUS members for their contributions to the FOCUS Release **v1.2** specification.

- Andrew Qu (Everest)
- Andrew Quigley (Northwestern Mutual)
- Beau Nelford (Anglepoint)
- David Dinh (Google)
- Erik Peterson (CloudZero)
- George Parker (Salesforce)
- Graham Murphy (Australian Retirement Trust)
- Janine Pickard-Green (MagicOrange Group Limited)
- Jason Wu (Amazon Web Services)
- Jesse Adams (Acuity)
- John Grubb (Platform.sh)
- Justin Grinstead (MongoDB)
- Marc Perreaut (Amadeus)
- Matt Cowser (FinOps Foundation)
- Nan Braun (Thavron Solutions)
- Ramkumar Narla (Adobe)
- Rob Martin (FinOps Foundation)
- Sanjna Srivatsa (CloudHealth)
- Steve Pearson (Capital One)
- Tim Wright (Google)
- Todd Blake (Datadog)

Steering Committee Members

Thanks to the following FOCUS Steering Committee members for their leadership on the FOCUS specification.

- Amit Kinha (Citi)
- Anne Johnston (Capital One)
- Christopher Harris (Datadog)
- Letian Feng (Amazon Web Services)
- Michael Flanakin (Microsoft)
- Mike Fuller (FinOps Foundation)
- Richard Steck (Adobe)
- Roy Wolman (Amazon Web Services) (term ended 2025/01/22)

- Sarah McMullin (Google)
- Tim O'Brien (Walmart)

Table of Contents

1. Introduction

- 1.1. Background and History
- 1.2. Intended Audience
- 1.3. Scope
- 1.4. Design Principles
- 1.5. Design Notes
- 1.6. Typographic Conventions
- 1.7. FOCUS Feature Level
- 1.8. Conformance Checkers and Validators

2. Supported Features

- 2.1. Account Structures
- 2.2. Billed Cost and Invoice Alignment
- 2.3. Charge Categorization
- 2.4. Commit Usage and Under Usage
- 2.5. Cost and Usage Attribution
- 2.6. Cost Comparison
- 2.7. Custom Columns
- 2.8. Data Granularity
- 2.9. Effective Cost Analysis
- 2.10. Location
- 2.11. Marketplace Purchases
- 2.12. Provider Services
- 2.13. Resource Usage
- 2.14. Schema Metadata
- 2.15. Service Categorization
- 2.16. Verification, Comparison, and Fluctuation Tracking of Unit Prices

3. Columns

- 3.1. Availability Zone
- 3.2. Billed Cost
- 3.3. Billing Account ID
- 3.4. Billing Account Name
- 3.5. Billing Account Type
- 3.6. Billing Currency
- 3.7. Billing Period End
- 3.8. Billing Period Start
- 3.9. Capacity Reservation ID
- 3.10. Capacity Reservation Status
- 3.11. Charge Category
- 3.12. Charge Class
- 3.13. Charge Description
- 3.14. Charge Frequency
- 3.15. Charge Period End
- 3.16. Charge Period Start
- 3.17. Commitment Discount Category
- 3.18. Commitment Discount ID
- 3.19. Commitment Discount Name
- 3.20. Commitment Discount Quantity
- 3.21. Commitment Discount Status
- 3.22. Commitment Discount Type
- 3.23. Commitment Discount Unit
- 3.24. Consumed Quantity
- 3.25. Consumed Unit
- 3.26. Contracted Cost
- 3.27. Contracted Unit Price
- 3.28. Effective Cost

- 3.29. Invoice ID
- 3.30. Invoice Issuer
- 3.31. List Cost
- 3.32. List Unit Price
- 3.33. Pricing Category
- 3.34. Pricing Currency
- 3.35. Pricing Currency Contracted Unit Price
- 3.36. Pricing Currency Effective Cost
- 3.37. Pricing Currency List Unit Price
- 3.38. Pricing Quantity
- 3.39. Pricing Unit
- 3.40. Provider
- 3.41. Publisher
- 3.42. Region ID
- 3.43. Region Name
- 3.44. Resource ID
- 3.45. Resource Name
- 3.46. Resource Type
- 3.47. Service Category
- 3.48. Service Name
- 3.49. Service Subcategory
- 3.50. SKU ID
- 3.51. SKU Meter
- 3.52. SKU Price Details
- 3.53. SKU Price ID
- 3.54. Sub Account ID
- 3.55. Sub Account Name
- 3.56. Sub Account Type
- 3.57. Tags

4. Attributes

- 4.1. Column Handling
- 4.2. Currency Format
- 4.3. Date/Time Format
- 4.4. Discount Handling
- 4.5. Key-Value Format
- 4.6. Null Handling
- 4.7. Numeric Format
- 4.8. String Handling
- 4.9. Unit Format

5. Metadata

- 5.1. Data Generator
- 5.2. Schema

6. Use Case Library

7. Glossary

8. Appendix

- 8.1. Examples: Commitment Discount Scenarios
- 8.2. Examples: Commitment Discount Flexibility
- 8.3. Examples: Metadata
- 8.4. Examples: SaaS
- 8.5. Grouping Constructs for Resources or Services
- 8.6. Origination of Cost Data

1. Introduction

This section is non-normative.

FOCUS is a standards development organization (SDO) formed to establish an open, consensus-driven standard for billing data. In the absence of a broadly adopted standard, infrastructure and service [providers](#) have relied on proprietary billing schemas and inconsistent terminology, making cost data difficult to normalize and act upon across environments. This lack of conformance has forced FinOps [practitioners](#) to develop best-effort custom normalization schemes for each provider, in order to perform essential FinOps capabilities such as chargeback, cost allocation, budgeting and forecasting.

The FOCUS Specification, developed by a global community of practitioners and vendors, defines a consistent, vendor-neutral approach to billing data. It is designed to improve interoperability between providers, reduce operational complexity, and enable greater transparency in cloud and SaaS cost management.

1.1. Background and History

This project is supported by the [FinOps Foundation](#). This work initially started under the Open Billing working group under the FinOps Foundation. The decision was made in Jan 2023 to begin to migrate the work to a newly formed project under the Linux Foundation called the FinOps Open Cost and Usage Specification (FOCUS) to better support the creation of a specification.

1.2. Intended Audience

This specification is designed to be used by three major groups:

- Billing data generators: Infrastructure and services *providers* that bill based on consumption, such as (but not limited to):
 - [Cloud Service Providers \(CSPs\)](#)
 - Software as a Service (SaaS) platforms
 - [Managed Service Providers \(MSPs\)](#)
 - Internal infrastructure and service platforms
- FinOps tool *providers*: Organizations that provide tools to assist with FinOps
- FinOps practitioners: Organizations and individuals consuming billing data for doing FinOps

1.3. Scope

The FOCUS working group will develop an open-source specification for billing data. The schema will define data [dimensions](#), [metrics](#), a set of attributes about billing data, and a common lexicon for describing billing data.

1.4. Design Principles

The following principles were considered while building the specification.

1.4.1. FOCUS is an iterative, living specification

- Incremental iterations of the specification released regularly will provide higher value to practitioners and allow feedback as the specification develops. The goal is not to get to a complete, finished specification in one pass.

1.4.2. Working backward with ease of adoption

- Aim to work backward from essential FinOps capabilities that practitioners need to perform to prioritize the dimensions, metrics and attributes of the cost and usage data that should be defined in the specification to fulfill that capability.
- Be FinOps scenario-driven. Define columns that answer scenario questions; don't look for scenarios to fit a column, each column must have a use case.
- Don't add dimensions or metrics to the specification just because it can be added.
- When defining the specification, consideration should be made to existing data already in the major providers' (AWS, GCP, Azure, OCI) datasets.
- As long as it solves the FinOps use case, there should be a preference to align with data that is already present in a majority of the major providers.
- Strive for simplicity. However, prioritize accuracy, clarity, and consistency.
- Strive to build columns that serve a single purpose, with clear and concise names and values.
- The specification should allow data to be presented free from jargon, using simple understandable terms, and be approachable.
- Naming and terms used should be carefully considered to avoid using terms for which the definition could be confused by the reader. If a term must be used which has either an unclear or multiple definitions, it should be clarified in the [glossary](#).
- The specification should provide all of the data elements necessary for the [Capabilities](#).

1.4.3. Provider-neutral approach by default

- While the schema, naming, terminology, and attributes of many providers are reviewed during development, this specification aims to be provider-neutral.
- Contributors must take care to ensure the specification examines how each decision relates to each of the major cloud providers and SaaS vendors, not favoring any single one.
- In some cases, the approach may closely resemble one or more provider's implementations, while in other cases, the approach might be new. In all cases, the FOCUS group (community composed of FinOps practitioners, Cloud and SaaS providers and FinOps vendors) will attempt to prioritize enabling FinOps [Capabilities](#) and alignment with the FinOps [Framework](#).

1.4.4. Extensibility

The FOCUS Specification is designed to support evolving FinOps needs across diverse billing models and provider types.

While the initial focus was on billing data from Cloud Service Providers (CSPs), version 1.2 introduces foundational support for Software as a Service (SaaS) platforms, including normative columns for pricing currencies, effective cost, and contracted pricing in non-monetary units such as credits or tokens.

The specification supports extensibility through structured naming conventions (e.g., x_ custom columns), conditional requirements, and a version-aware schema approach.

Future versions of FOCUS will consider including additional FinOps capabilities such as forecasting, exchange rate modeling, and anomaly detection, while continuing to support a broader range of billing and cost datasets — including internal infrastructure platforms and marketplace offerings.

1.5. Design Notes

1.5.1. Optimize for data analysis

- Optimize columns for data analysis at scale and avoid the requirement of splitting or parsing values.
- Avoid complex JSON structures when an alternative columnar structure is possible.
- Facilitate the inclusion of data necessary for a system of record for cost and usage data to consume.

1.5.2. Consistency helps with clarity

- Where possible, use consistent names that will naturally create associations between related columns in the specification.
- Column naming must strictly follow the [column handling](#) requirements.
- Use established standards (e.g., ISO8601 for dates, ISO4217 for currency).

1.6. Typographic Conventions

The keywords "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this specification are to be interpreted as described in [BCP14](#) [[RFC2119](#)][[RFC8174](#)] when, and only when, they appear in all capitals, as shown here.

1.7. FOCUS Feature Level

Under each column defined in the FOCUS specification, there exists a 'Feature level' designation that describes the column as 'Mandatory', 'Conditional', or 'Optional'. Feature level is designated based on the following criteria described in the normative requirements in each column definition:

- If the existence of a column is described with MUST with no conditions of when it applies, then the feature level is designated as 'Mandatory'.
- If the existence of a column is described as MUST with conditions of when it applies, then the feature level is designated as 'Conditional'.
- If the existence of a column is described as RECOMMENDED, then the feature level is designated as 'Recommended'.
- If the existence of a column is described as MAY, then the feature level is designated as 'Optional'.

1.8. Conformance Checkers and Validators

There are no current resources available to test for specification conformance or validators to run on sample data. When one becomes available, this section of the specification will be updated with details.

2. Supported Features

The FOCUS specification is designed to meet the needs of FinOps practitioners in numerous scenarios. The following section contains features supported by the FOCUS specification. This list does not represent all possible combinations or use of FOCUS data but does represent core capabilities that the FOCUS specification supports.

2.1. Account Structures

2.1.1. Description

Different providers have different account constructs that FinOps practitioners use for allocation, reporting, and more. Organizations may have one or many accounts within one or more providers and FinOps practitioners may need to review the cost broken down by each account. FOCUS has two types of accounts: a billing account and a sub account.

A billing account is the account where invoices are generated. Each billing account can have one or more sub accounts, which can be used for deploying and managing resources and services. Billing and sub accounts are often used to facilitate allocation strategies and FinOps practitioners must be able to break costs down by billing and sub account to facilitate FinOps scenarios like chargeback and budgeting.

2.1.2. Directly Dependent Columns

- BilledCost
- BillingAccountId
- BillingAccountName
- BillingAccountType
- SubAccountId
- SubAccountName
- SubAccountType

2.1.3. Supporting Columns

- InvoiceId

2.1.4. Example SQL Query

SELECT

BillingAccountId,
BillingAccountName,
BillingAccountType,
SubAccountId,
SubAccountName,
SubAccountType,
SUM(BilledCost)

FROM focus_data_table

WHERE BillingPeriodStart >= ? **AND** BillingPeriodEnd < ?

GROUP BY

BillingAccountId,
SubAccountId

2.1.5. Introduced (Version)

0.5

2.2. Billed Cost and Invoice Alignment

2.2.1. Description

FOCUS data should be consistent with the costs indicated on payable invoices. This is relevant to the total cost of the invoice, as well as the period of time the invoice covers.

2.2.2. Directly Dependent Columns

- BilledCost
- BillingCurrency
- BillingPeriodEnd
- BillingPeriodStart
- InvoiceId

2.2.3. Supporting Columns

- ProviderName
- ServiceName

2.2.4. Example SQL Query

SELECT

BillingPeriodStart,
BillingPeriodEnd,
InvoiceId,
SUM(BilledCost)

FROM focus_data_table

GROUP BY

BillingPeriodStart,
BillingPeriodEnd,
InvoiceId

2.2.5. Introduced (Version)

0.5

2.3. Charge Categorization

2.3.1. Description

FOCUS supports the categorization of charges including purchases, usage, tax, credits and adjustments. It includes classification on frequency. It includes classification on correction vs normal entries

2.3.2. Directly Dependent Columns

- ChargeCategory
- ChargeClass
- ChargeFrequency

2.3.3. Supporting Columns

- BilledCost
- BillingAccountId
- BillingPeriodEnd
- BillingPeriodStart
- CommitmentDiscountId
- CommitmentDiscountType
- ProviderName
- ServiceCategory

2.3.4. Example SQL Query

2.3.4.1. Report on Commitment Discount Purchases

SELECT

MIN(ChargePeriodStart) **AS** ChargePeriodStart,
MAX(ChargePeriodEnd) **AS** ChargePeriodEnd,
ProviderName,
BillingAccountId,
CommitmentDiscountId,
CommitmentDiscountType,
CommitmentDiscountUnit,
CommitmentDiscountQuantity,
ChargeFrequency,
SUM(BilledCost) **AS** TotalBilledCost

FROM focus_data_table

WHERE ChargePeriodStart >= ? **AND** ChargePeriodEnd < ?

AND ChargeCategory = 'Purchase'

AND CommitmentDiscountId **IS NOT NULL**

GROUP BY

ProviderName,
BillingAccountId,
CommitmentDiscountId,
CommitmentDiscountType,
CommitmentDiscountUnit,
CommitmentDiscountQuantity,
ChargeFrequency

2.3.4.2. Report on Corrections**SELECT**

ProviderName,
BillingAccountId,
ChargeCategory,
ServiceCategory,
ServiceName,
SUM(BilledCost) **AS** TotalBilledCost

FROM focus_data_table

WHERE BillingPeriodStart >= ? **AND** BillingPeriodEnd < ?

AND ChargeClass = 'Correction'

GROUP BY

ProviderName,
BillingAccountId,
ChargeCategory,
ServiceCategory,
ServiceName

2.3.4.3. Report Recurring Charges

SELECT

```
BillingPeriodStart,  
CommitmentDiscountId,  
CommitmentDiscountName,  
CommitmentDiscountType,  
ChargeFrequency,  
SUM(BilledCost) AS TotalBilledCost
```

FROM focus_data_table

WHERE BillingPeriodStart >= ? **AND** BillingPeriodStart < ?

AND ChargeFrequency = 'Recurring'

AND CommitmentDiscountId **IS NOT NULL**

GROUP BY

```
BillingPeriodStart,  
CommitmentDiscountId,  
CommitmentDiscountName,  
CommitmentDiscountType,  
ChargeFrequency
```

2.3.5. Introduced (Version)

1.0

2.4. Commit Usage and Under Usage

2.4.1. Description

FOCUS supports the tracking of commitment discounts usage and under usage, which can come in the form of commitment discounts or capacity reservations.

2.4.2. Directly Dependent Columns

- CommitmentDiscountID
- CommitmentDiscountStatus
- CommitmentDiscountType
- CapacityReservationID
- CapacityReservationStatus
- CapacityReservationType

2.4.3. Supporting Columns

- BilledCost
- ChargePeriodStart
- ChargePeriodEnd
- EffectiveCost
- ServiceCategory

2.4.4. Example SQL Query for Commitment Discounts

```

SELECT
  ProviderName,
  BillingAccountId,
  CommitmentDiscountId,
  CommitmentDiscountType,
  CommitmentDiscountStatus,
  SUM(BilledCost) AS TotalBilledCost,
  SUM(EffectiveCost) AS TotalEffectiveCost
FROM focus_data_table
WHERE ChargePeriodStart >= ? AND ChargePeriodEnd < ?
  AND CommitmentDiscountStatus = 'Unused'
GROUP BY
  ProviderName,
  BillingAccountId,
  CommitmentDiscountId,
  CommitmentDiscountType

```

2.4.5. Example SQL Query for Capacity Reservations

```

SELECT
  ProviderName,
  BillingAccountId,
  CapacityReservationId,
  CapacityReservationStatus,
  SUM(BilledCost) AS TotalBilledCost,
  SUM(EffectiveCost) AS TotalEffectiveCost
FROM focus_data_table
WHERE ChargePeriodStart >= ? AND ChargePeriodEnd < ?
  AND CapacityReservationStatus = 'Unused'
GROUP BY
  ProviderName,
  BillingAccountId,
  CapacityReservationId,
  CapacityReservationStatus

```

2.4.6. Introduced (Version)

1.0

2.5. Cost and Usage Attribution

2.5.1. Description

Many providers have features that allow Finops practitioners to enrich cost and usage data with metadata, that is addition to provider defined data, in order to analyze Finops data using organizational, deployment, or other structures. These features may take the form of directly applied metadata or inherited metadata. FOCUS facilitates the inclusion of this metadata at a row level.

2.5.2. Directly Dependent Columns

- Tags

2.5.3. Supporting Columns

- BilledCost
- ConsumedQuantity
- ConsumedUnit
- EffectiveCost

2.5.4. Example SQL Query

SELECT

```
tags,  
ConsumedUnit,  
SUM(BilledCost),  
SUM(EffectiveCost),  
SUM(ConsumedQuantity)
```

FROM focus_data_table

WHERE BillingPeriodStart >= ? **AND** BillingPeriodEnd < ?

GROUP BY

```
tags,  
ConsumedUnit
```

2.5.5. Introduced (Version)

1.0

2.6. Cost Comparison

2.6.1. Description

FOCUS supports the comparison of cost columns in order to identify savings, amortization, or other constructs.

2.6.2. Directly Dependent Columns

- BilledCost
- ContractedCost
- EffectiveCost
- ListCost

2.6.3. Supporting Columns

- BillingAccountId
- BillingAccountName
- BillingCurrency

- BillingPeriodEnd
- BillingPeriodStart
- ChargePeriodEnd
- ChargePeriodStart
- ServiceName

2.6.4. Example SQL Query

```

WITH AggregatedData AS (
  SELECT
    ProviderName,
    BillingAccountId,
    BillingAccountName,
    BillingCurrency,
    ServiceName,
    SUM(EffectiveCost) AS TotalEffectiveCost,
    SUM(BilledCost) AS TotalBilledCost,
    SUM(CASE
      WHEN ChargeCategory = 'Usage' AND BilledCost = 0 AND EffectiveCost != 0
      THEN 0
      ELSE ContractedCost
    END) AS TotalContractedCost,
    SUM(CASE
      WHEN ChargeCategory = 'Usage' AND BilledCost = 0 AND EffectiveCost != 0
      THEN 0
      ELSE ListCost
    END) AS TotalListCost
  FROM focus_data_table
  WHERE BillingPeriodStart >= ?
    AND BillingPeriodEnd < ?
    AND ChargeClass IS NULL
  GROUP BY
    ProviderName,
    BillingAccountId,
    BillingAccountName,
    BillingCurrency,
    ServiceName
)
SELECT ProviderName,
  BillingAccountId,
  BillingAccountName,
  BillingCurrency,
  ServiceName,
  TotalEffectiveCost,
  TotalBilledCost,
  TotalListCost,
  1 - (TotalContractedCost / NULLIF(TotalListCost, 0)) * 100 AS ContractedDiscount
  1 - (TotalEffectiveCost / NULLIF(TotalListCost, 0)) * 100 AS EffectiveDiscount
FROM AggregatedData

```

2.6.5. Introduced (Version)

0.5

2.7. Custom Columns

2.7.1. Description

FOCUS supports the inclusion of custom columns to facilitate reporting capability that is not covered by the columns included in the specification.

2.7.2. Directly Dependent Columns

- x_CustomColumn

2.7.3. Example SQL Query

```
SELECT
    BillingPeriodStart,
    x_CustomColumn,
    SUM(BilledCost) AS TotalBilledCost,
FROM focus_data_table
WHERE ServiceName = ?
    AND BillingPeriodStart >= ? AND BillingPeriodStart < ?
GROUP BY
    BillingPeriodStart,
    x_CustomColumn
ORDER BY MonthlyCost DESC
```

2.7.4. Introduced (Version)

0.5

2.8. Data Granularity

2.8.1. Description

FOCUS supports multiple levels of cost and usage data granularity. This includes the ability to report on a daily, hourly, or other time period basis. FOCUS also supports the ability for cost and usage data to be provided for high granularity scenarios, such as down to the individual resources. It also supports high level granularity cost and usage data, such as account level, or service level charges.

2.8.2. Directly Dependent Columns

- ResourceId
- ResourceName
- ChargePeriodEnd
- ChargePeriodStart

2.8.3. Supporting Columns

- BilledCost
- ConsumedQuantity
- ConsumedUnit
- EffectiveCost
- ListCost
- PricingCurrency
- PricingUnit

2.8.4. Example SQL Query

SELECT

ChargePeriodStart,
ChargePeriodEnd,
ResourceId,
SUM(EffectiveCost)

FROM focus_data_table

Group by

ChargePeriodStart,
ChargePeriodEnd,
ResourceId

2.8.5. Introduced (Version)

0.5

2.9. Effective Cost Analysis

2.9.1. Description

FOCUS enables practitioners to analyze costs without having to distribute upfront fees and discounts, taking discounts and the amortization of upfront fees paid for services into account. The EffectiveCost column represents cost after negotiated discounts, commitment discounts, and the applicable portion of relevant, prepaid purchases (one-time or recurring) that covered this charge. EffectiveCost is commonly utilized to track and analyze spending trends.

2.9.2. Directly Dependent Columns

- EffectiveCost

2.9.3. Supporting Columns

- BillingPeriodEnd
- BillingPeriodStart
- ChargeCategory
- ChargePeriodEnd
- ChargePeriodStart

- ConsumedQuantity
- ConsumedUnit
- PricingQuantity
- ProviderName
- RegionName
- ServiceName

2.9.4. Example SQL Query

SELECT

```
ProviderName,
BillingPeriodStart,
BillingPeriodEnd,
ServiceCategory,
ServiceName,
RegionId,
RegionName,
PricingUnit,
SUM(EffectiveCost) AS TotalEffectiveCost,
SUM(PricingQuantity) AS TotalPricingQuantity
```

FROM focus_data_table

WHERE BillingPeriodStart >= ? **AND** BillingPeriodEnd <= ?

GROUP BY

```
ProviderName,
BillingPeriodStart,
BillingPeriodEnd,
ServiceCategory,
ServiceName,
RegionId,
RegionName,
PricingUnit
```

2.9.5. Introduced (Version)

0.5

2.10. Location

2.10.1. Description

FOCUS provides structured location data through region and availability zone information. By documenting geographic deployment locations, practitioners can organize and analyze costs based on where resources and services are deployed. This standardized location data helps practitioners understand the geographical distribution of infrastructure across providers.

2.10.2. Directly Dependent Columns

- AvailabilityZone
- RegionId
- RegionName

2.10.3. Supporting Columns

- BilledCost
- ChargePeriodEnd
- ChargePeriodStart

2.10.4. Example SQL Query

SELECT

RegionId,
RegionName,
AvailabilityZone,
SUM(BilledCost) AS TotalBilledCost

FROM focus_data_table

WHERE ChargePeriodStart >= ? AND ChargePeriodEnd <= ?

GROUP BY

RegionId,
RegionName,
AvailabilityZone

2.10.5. Introduced (Version)

1.0

2.11. Marketplace Purchases

2.11.1. Description

FOCUS supports the analysis of cost and usage data for marketplace purchases and their associated costs. It also supports the reporting of EffectiveCost for usage from the provider.

2.11.2. Directly Dependent Columns

- InvoiceIssuer
- Provider

2.11.3. Supporting Columns

- BilledCost
- EffectiveCost

2.11.4. Example SQL Query on a CSP Marketplace FOCUS Dataset

```

SELECT
  Provider,
  InvoiceIssuer,
  BillingPeriodStart,
  BillingPeriodEnd,
  SUM(BilledCost) AS TotalBilledCost
FROM focus_data_table
WHERE Provider = '<Example SaaS Provider>'
  AND InvoiceIssuer = '<Example CSP Marketplace>'
GROUP BY
  Provider,
  InvoiceIssuer,
  BillingPeriodStart,
  BillingPeriodEnd

```

2.11.5. Example SQL Query on a Provider FOCUS Dataset

```

SELECT
  ChargePeriodStart,
  ChargePeriodEnd,
  ResourceId,
  SUM(EffectiveCost) AS TotalEffectiveCost
FROM focus_data_table
WHERE InvoiceIssuer = '<Example CSP Marketplace>'
GROUP BY
  ChargePeriodStart,
  ChargePeriodEnd,
  ResourceId

```

2.11.6. Introduced (Version)

1.0

2.12. Provider Services

2.12.1. Description

FOCUS supports providers specifying the services and product offerings that they provide their customers that align with the names practitioners are familiar with. This empowers practitioners to analyze cost by service, report service costs by subaccount, forecast based on historical trends by service, and verify accuracy of services charged across providers.

2.12.2. Directly Dependent Columns

- ServiceCategory
- ServiceName
- ServiceSubcategory

2.12.3. Supporting Columns

- Provider
- Skuld

2.12.4. Example SQL Query

SELECT

```
BillingPeriodStart,  
ProviderName,  
SubAccountId,  
SubAccountName,  
ServiceName,  
SUM(BilledCost) AS TotalBilledCost,  
SUM(EffectiveCost) AS TotalEffectiveCost
```

FROM focus_data_table

WHERE ServiceName = ?

AND BillingPeriodStart >= ? **AND** BillingPeriodStart < ?

GROUP BY

```
BillingPeriodStart,  
ProviderName,  
SubAccountId,  
SubAccountName,  
ServiceName
```

ORDER BY MonthlyCost **DESC**

2.12.5. Introduced (Version)

0.5

2.13. Resource Usage

2.13.1. Description

FOCUS enables tracking of resource consumption by providing information about which resources were used, in what quantities, and with what units of measure.

2.13.2. Directly Dependent Columns

- ConsumedQuantity
- ConsumedUnit
- ResourceId
- Skuld

2.13.3. Supporting Columns

- ChargeCategory
- ChargePeriodEnd

- ChargePeriodStart
- ProviderName
- ServiceName

2.13.4. Example SQL Query

SELECT

```

ProviderName,
ServiceName,
ResourceId,
Skuld,
ConsumedUnit,
SUM(ConsumedQuantity) AS TotalQuantity
FROM focus_data_table
WHERE ChargeCategory='Usage'
AND ChargePeriodStart >= ? AND ChargePeriodEnd <= ?
GROUP BY
ProviderName,
ServiceName,
ResourceId,
Skuld,
ConsumedUnit

```

2.13.5. Introduced (Version)

1.0

2.14. Schema Metadata

2.14.1. Description

FOCUS' schema metadata supports communication of important attributes about the data, facilitating notifications about changing structure and database table creation between provider and consumer. This includes column names, data types, and any other relevant information about the data schema. It also includes information as to the version of FOCUS and Data Generator versioning that the data uses.

2.14.2. Applicable Metadata

- Schema
 - Column Definition

2.14.3. Introduced (Version)

1.1

2.15. Service Categorization

2.15.1. Description

FOCUS provides a structure for categorizing services based on their core functions. By classifying services into high-level categories and more granular subcategories, practitioners can organize costs according to functional areas. This standardized categorization provides data that practitioners can use in their cost management processes and decision making.

2.15.2. Directly Dependent Columns

- ServiceCategory
- ServiceName
- ServiceSubcategory

2.15.3. Supporting Columns

- BilledCost
- BillingCurrency
- BillingPeriodEnd
- BillingPeriodStart
- ProviderName

2.15.4. Example SQL Query

SELECT

BillingPeriodStart,
BillingPeriodEnd,
ProviderName,
ServiceCategory,
ServiceSubcategory,
ServiceName,
BillingCurrency,
SUM(BilledCost) **AS** TotalBilledCost

FROM focus_data_table

WHERE BillingPeriodStart >= ? **and** BillingPeriodEnd < ?

GROUP BY

BillingPeriodStart,
BillingPeriodEnd,
ProviderName,
ServiceCategory,
ServiceSubcategory,
ServiceName,
BillingCurrency

2.15.5. Introduced (Version)

1.1

2.16. Verification, Comparison, and Fluctuation Tracking of Unit Prices

2.16.1. Description

When a provider supports unit pricing concepts, FOCUS allows practitioners to:

- Verify that the correct List Unit Prices and Contracted Unit Prices are applied.
- Compare applied Contracted Unit Prices across different billing accounts and with applied List Unit Prices at specific points in time.
- Track fluctuations in unit prices over time.

2.16.2. Directly Dependent Columns

- ContractedUnitPrice
- ListUnitPrice
- Skuld
- SkuPriceDetails
- SkuPricId

2.16.3. Supporting Columns

- BillingCurrency
- BillingPeriodId
- ChargePeriodEnd
- ChargePeriodStart

2.16.4. Example SQL Query

SELECT DISTINCT

Skuld,
SkuPricId,
SkuPriceDetails,
BillingPeriodId,
ChargePeriodStart,
ChargePeriodEnd,
BillingCurrency,
ListUnitPrice,
ContractedUnitPrice

FROM focus_data_table

WHERE

SkuPricId = ?

AND ChargePeriodStart >= ?

AND ChargePeriodEnd < ?

2.16.5. Introduced (Version)

1.0

3. Columns

The FOCUS specification defines a group of columns that provide qualitative values (such as dates, resource, and provider information) categorized as "dimensions" and quantitative values (numeric values) categorized as "metrics" that can be used for performing various [FinOps capabilities](#). Metrics are commonly used for aggregations (sum, multiplication, averaging etc.) and statistical operations within the dataset. Dimensions are commonly used to categorize, filter, and reveal details in your data when combined with metrics. The Columns are presented in Alphabetical order.

3.1. Availability Zone

An [availability zone](#) is a provider-assigned identifier for a physically separated and isolated area within a Region that provides high availability and fault tolerance. Availability Zone is commonly used for scenarios like analyzing cross-zone data transfer usage and the corresponding cost based on where [resources](#) are deployed.

The AvailabilityZone column adheres to the following requirements:

- AvailabilityZone is RECOMMENDED to be present in a [FOCUS dataset](#) when the provider supports deploying resources or services within an *availability zone*.
- AvailabilityZone MUST be of type String.
- AvailabilityZone MUST conform to [StringHandling](#) requirements.
- AvailabilityZone MUST be null when a [charge](#) is not specific to an *availability zone*.

3.1.1. Column ID

AvailabilityZone

3.1.2. Display Name

Availability Zone

3.1.3. Description

A provider-assigned identifier for a physically separated and isolated area within a Region that provides high availability and fault tolerance.

3.1.4. Content constraints

Constraint	Value
Column type	Dimension
Feature level	Recommended
Allows nulls	True
Data type	String

Constraint	Value
Value format	<not specified>

3.1.5. Introduced (version)

0.5

3.2. Billed Cost

The [billed cost](#) represents a [charge](#) serving as the basis for invoicing, inclusive of the impacts of all reduced rates and discounts while excluding the [amortization](#) of relevant purchases (one-time or recurring) paid to cover future eligible [charges](#). This cost is denominated in the [Billing Currency](#). The Billed Cost is commonly used to perform FinOps capabilities that require cash-basis accounting such as cost allocation, budgeting, and invoice reconciliation.

The BilledCost column adheres to the following requirements:

- BilledCost MUST be present in a [FOCUS dataset](#).
- BilledCost MUST be of type Decimal.
- BilledCost MUST conform to [NumericFormat](#) requirements.
- BilledCost MUST NOT be null.
- BilledCost MUST be a valid decimal value.
- BilledCost MUST be 0 for [charges](#) where payments are received by a third party (e.g., marketplace transactions).
- BilledCost MUST be denominated in the BillingCurrency.
- The sum of the BilledCost for a given [InvoiceId](#) MUST match the sum of the payable amount provided in the corresponding invoice with the same id generated by the [InvoiceIssuer](#).

3.2.1. Column ID

BilledCost

3.2.2. Display Name

Billed Cost

3.2.3. Description

A [charge](#) serving as the basis for invoicing, inclusive of all reduced rates and discounts while excluding the [amortization](#) of upfront [charges](#) (one-time or recurring).

3.2.4. Content constraints

Constraint	Value
Column type	Metric
Feature level	Mandatory
Allows nulls	False

Constraint	Value
Data type	Decimal
Value format	Numeric Format
Number range	Any valid decimal value

3.2.5. Introduced (version)

0.5

3.3. Billing Account ID

A Billing Account ID is a provider-assigned identifier for a [billing account](#). *Billing accounts* are commonly used for scenarios like grouping based on organizational constructs, invoice reconciliation and cost allocation strategies.

The BillingAccountId column adheres to the following requirements:

- BillingAccountId MUST be present in a [FOCUS dataset](#).
- BillingAccountId MUST be of type String.
- BillingAccountId MUST conform to [StringHandling](#) requirements.
- BillingAccountId MUST NOT be null.
- BillingAccountId MUST be a unique identifier within a provider.
- BillingAccountId SHOULD be a fully-qualified identifier.

See [Appendix: Grouping constructs for resources or services](#) for details and examples of the different grouping constructs supported by FOCUS.

3.3.1. Column ID

BillingAccountId

3.3.2. Display Name

Billing Account ID

3.3.3. Description

The identifier assigned to a *billing account* by the provider.

3.3.4. Content constraints

Constraint	Value
Column type	Dimension
Feature level	Mandatory
Allows nulls	False

Constraint	Value
Data type	String
Value format	<not specified>

3.3.5. Introduced (version)

0.5

3.4. Billing Account Name

A Billing Account Name is a display name assigned to a [billing account](#). *Billing accounts* are commonly used for scenarios like grouping based on organizational constructs, invoice reconciliation and cost allocation strategies.

The BillingAccountName column adheres to the following requirements:

- BillingAccountName MUST be present in a [FOCUS dataset](#).
- BillingAccountName MUST be of type String.
- BillingAccountName MUST conform to [StringHandling](#) requirements.
- BillingAccountName MUST NOT be null when the provider supports assigning a display name for the *billing account*.

See [Appendix: Grouping constructs for resources or services](#) for details and examples of the different grouping constructs supported by FOCUS.

3.4.1. Column ID

BillingAccountName

3.4.2. Display Name

Billing Account Name

3.4.3. Description

The display name assigned to a *billing account*.

3.4.4. Content constraints

Constraint	Value
Column type	Dimension
Feature level	Mandatory
Allows nulls	True
Data type	String
Value format	<not specified>

3.4.5. Introduced (version)

0.5

3.5. Billing Account Type

Billing Account Type is a provider-assigned name to identify the type of [billing account](#). Billing Account Type is a readable display name and not a code. Billing Account Type is commonly used for scenarios like mapping FOCUS and provider constructs, summarizing costs across providers, or invoicing and chargeback.

The BillingAccountType column adheres to the following requirements:

- BillingAccountType MUST be present in a [FOCUS dataset](#) when the provider supports more than one possible BillingAccountType value.
- BillingAccountType MUST be of type String.
- BillingAccountType MUST conform to [StringHandling](#) requirements.
- BillingAccountType nullability is defined as follows:
 - BillingAccountType MUST be null when [BillingAccountId](#) is null.
 - BillingAccountType MUST NOT be null when BillingAccountId is not null.
- BillingAccountType MUST be a consistent, readable display value.

3.5.1. Column ID

BillingAccountType

3.5.2. Display Name

Billing Account Type

3.5.3. Description

A provider-assigned name to identify the type of *billing account*.

3.5.4. Content Constraints

Constraint	Value
Column type	Dimension
Feature level	Conditional
Allows nulls	False
Data type	String
Value format	<not specified>

3.5.5. Introduced (version)

3.6. Billing Currency

[Billing currency](#) is an identifier that represents the currency that a [charge](#) for [resources](#) or [services](#) was billed in. Billing Currency is commonly used in scenarios where costs need to be grouped or aggregated.

The BillingCurrency column adheres to the following requirements:

- BillingCurrency MUST be present in a [FOCUS dataset](#).
- BillingCurrency MUST be of type String.
- BillingCurrency MUST conform to [StringHandling](#) requirements.
- BillingCurrency MUST conform to [CurrencyFormat](#) requirements.
- BillingCurrency MUST NOT be null.
- BillingCurrency MUST match the currency used in the invoice generated by the invoice issuer.
- BillingCurrency MUST be expressed in [national currency](#) (e.g., USD, EUR).

3.6.1. Column ID

BillingCurrency

3.6.2. Display Name

Billing Currency

3.6.3. Description

Represents the currency that a *charge* was billed in.

3.6.4. Content Constraints

Constraint	Value
Column type	Dimension
Feature level	Mandatory
Allows nulls	False
Data type	String
Value format	Currency Format

3.6.5. Introduced (version)

0.5

3.7. Billing Period End

Billing Period End represents the [exclusive end bound](#) of a [billing period](#). For example, a time period where [Billing Period Start](#) is '2024-01-01T00:00:00Z' and Billing Period End is '2024-02-01T00:00:00Z' includes [charges](#) for January since Billing Period Start represents the [inclusive start bound](#), but does not include [charges](#) for February since Billing Period End represents the [exclusive end bound](#).

The BillingPeriodEnd column adheres to the following requirements:

- BillingPeriodEnd MUST be present in a [FOCUS dataset](#).
- BillingPeriodEnd MUST be of type Date/Time.
- BillingPeriodEnd MUST conform to [DateTimeFormat](#) requirements.
- BillingPeriodEnd MUST NOT be null.
- BillingPeriodEnd MUST be the [exclusive end bound](#) of the [billing period](#).

3.7.1. Column ID

BillingPeriodEnd

3.7.2. Display Name

Billing Period End

3.7.3. Description

The [exclusive end bound](#) of a [billing period](#).

3.7.4. Content Constraints

Constraint	Value
Column type	Dimension
Feature level	Mandatory
Allows nulls	False
Data type	Date/Time
Value format	Date/Time Format

3.7.5. Introduced (version)

0.5

3.8. Billing Period Start

Billing Period Start represents the [inclusive start bound](#) of a [billing period](#). For example, a time period where Billing Period Start is '2024-01-01T00:00:00Z' and [Billing Period End](#) is '2024-02-01T00:00:00Z' includes [charges](#) for January since Billing Period Start represents the [inclusive start bound](#), but does not include [charges](#) for February since BillingPeriodEnd represents the [exclusive end bound](#).

The BillingPeriodStart column adheres to the following requirements:

- BillingPeriodStart MUST be present in a [FOCUS dataset](#).
- BillingPeriodStart MUST be of type Date/Time.
- BillingPeriodStart MUST conform to [DateTimeFormat](#) requirements.
- BillingPeriodStart MUST NOT be null.
- BillingPeriodStart MUST be the *inclusive start bound* of the *billing period*.

3.8.1. Column ID

BillingPeriodStart

3.8.2. Display Name

Billing Period Start

3.8.3. Description

The *inclusive start bound* of a *billing period*.

3.8.4. Content Constraints

Constraint	Value
Column type	Dimension
Feature level	Mandatory
Allows nulls	False
Data type	Date/Time
Value format	Date/Time Format

3.8.5. Introduced (version)

0.5

3.9. Capacity Reservation ID

A Capacity Reservation ID is the identifier assigned to a [capacity reservation](#) by the provider. Capacity Reservation ID is commonly used for scenarios to allocate [charges](#) for capacity reservation usage.

The CapacityReservationId column adheres to the following requirements:

- CapacityReservationId MUST be present in a [FOCUS dataset](#) when the provider supports *capacity reservations*.
- CapacityReservationId MUST be of type String.
- CapacityReservationId MUST conform to [StringHandling](#) requirements.
- CapacityReservationId nullability is defined as follows:
 - CapacityReservationId MUST be null when a *charge* is not related to a *capacity reservation*.

- CapacityReservationId MUST NOT be null when a *charge* represents the unused portion of a *capacity reservation*.
- CapacityReservationId SHOULD NOT be null when a *charge* is related to a capacity reservation.
- When CapacityReservationId is not null, CapacityReservationId adheres to the following additional requirements:
 - CapacityReservationId MUST be a unique identifier within the provider.
 - CapacityReservationId SHOULD be a fully-qualified identifier.

3.9.1. Column ID

CapacityReservationId

3.9.2. Display Name

Capacity Reservation ID

3.9.3. Description

The identifier assigned to a *capacity reservation* by the provider.

3.9.4. Content constraints

Constraint	Value
Column type	Dimension
Feature level	Conditional
Allows nulls	True
Data type	String
Value format	<not specified>

3.9.5. Introduced (version)

1.1

3.10. Capacity Reservation Status

Capacity Reservation Status indicates whether the [charge](#) represents either the consumption of the [capacity reservation](#) identified in the CapacityReservationId column or when the *capacity reservation* is unused.

The CapacityReservationStatus column adheres to the following requirements:

- CapacityReservationStatus MUST be present in a [FOCUS dataset](#) when the provider supports *capacity reservations*.
- CapacityReservationStatus MUST be of type String.
- CapacityReservationStatus nullability is defined as follows:
 - CapacityReservationStatus MUST be null when CapacityReservationId is null.

- CapacityReservationStatus MUST NOT be null when CapacityReservationId is not null and [ChargeCategory](#) is "Usage".
- When CapacityReservationStatus is not null, CapacityReservationStatus adheres to the following additional requirements:
 - CapacityReservationStatus MUST be one of the allowed values.
 - CapacityReservationStatus MUST be "Unused" when the *charge* represents the unused portion of a *capacity reservation*.
 - CapacityReservationStatus MUST be "Used" when the *charge* represents the used portion of a *capacity reservation*.

3.10.1. Column ID

CapacityReservationStatus

3.10.2. Display Name

Capacity Reservation Status

3.10.3. Description

Indicates whether the *charge* represents either the consumption of a *capacity reservation* or when a *capacity reservation* is unused.

3.10.4. Content constraints

Constraint	Value
Column type	Dimension
Feature level	Conditional
Allows nulls	True
Data type	String
Value format	Allowed Values

Allowed values:

Value	Description
Used	<i>Charges</i> that utilized a specific amount of a <i>capacity reservation</i> .
Unused	<i>Charges</i> that represent the unused portion of a <i>capacity reservation</i> .

3.10.5. Introduced (version)

1.1

3.11. Charge Category

Charge Category represents the highest-level classification of a [charge](#) based on the nature of how it is billed. Charge Category is commonly used to identify and distinguish between types of [charges](#) that may require different handling.

The ChargeCategory column adheres to the following requirements:

- ChargeCategory MUST be present in a [FOCUS dataset](#).
- ChargeCategory MUST be of type String.
- ChargeCategory MUST NOT be null.
- ChargeCategory MUST be one of the allowed values.

3.11.1. Column ID

ChargeCategory

3.11.2. Display Name

Charge Category

3.11.3. Description

Represents the highest-level classification of a *charge* based on the nature of how it is billed.

3.11.4. Content Constraints

Constraint	Value
Column type	Dimension
Feature level	Mandatory
Allows nulls	False
Data type	String
Value format	Allowed values

Allowed values:

Value	Description
Usage	Positive or negative <i>charges</i> based on the quantity of a service or resource that was consumed over a given period of time including refunds.
Purchase	Positive or negative <i>charges</i> for the acquisition of a service or resource bought upfront or on a recurring basis including refunds.
Tax	Positive or negative applicable taxes that are levied by the relevant authorities including refunds. Tax <i>charges</i> may vary depending on factors such as the location, jurisdiction, and local or federal regulations.
Credit	Positive or negative <i>charges</i> granted by the provider for various scenarios e.g promotional credits or corrections to promotional credits.
Adjustment	Positive or negative <i>charges</i> the provider applies that do not fall into other category values.

3.11.5. Introduced (version)

0.5

3.12. Charge Class

Charge Class indicates whether the [row](#) represents a correction to a previously invoiced [billing period](#). Charge Class is commonly used to differentiate [*corrections](#) from regularly incurred [charges](#).

The ChargeClass column adheres to the following requirements:

- ChargeClass MUST be present in a [FOCUS dataset](#).
- ChargeClass MUST be of type String.
- ChargeClass nullability is defined as follows:
 - ChargeClass MUST be null when the *row* does not represent a correction or when it represents a correction within the current *billing period*.
 - ChargeClass MUST NOT be null when the *row* represents a correction to a previously invoiced *billing period*.
- ChargeClass MUST be "Correction" when ChargeClass is not null.

3.12.1. Column ID

ChargeClass

3.12.2. Display Name

Charge Class

3.12.3. Description

Indicates whether the *row* represents a correction to a previously invoiced *billing period*.

3.12.4. Content Constraints

Constraint	Value
Column type	Dimension
Feature level	Mandatory
Allows nulls	True
Data type	String
Value format	Allowed values

Allowed values:

Value	Description
Correction	Correction to a previously invoiced <i>billing period</i> (e.g., refunds and credit modifications).

3.12.5. Introduced (version)

1.0

3.13. Charge Description

A Charge Description provides a high-level context of a [row](#) without requiring additional discovery. This column is a self-contained summary of the [charge's](#) purpose and price. It typically covers a select group of corresponding details across a billing dataset or provides information not otherwise available.

The ChargeDescription column adheres to the following requirements:

- ChargeDescription MUST be present in a [FOCUS dataset](#).
- ChargeDescription MUST be of type String.
- ChargeDescription MUST conform to [StringHandling](#) requirements.
- ChargeDescription SHOULD NOT be null.
- ChargeDescription maximum length SHOULD be provided in the corresponding FOCUS Metadata Schema.

3.13.1. Column ID

ChargeDescription

3.13.2. Display Name

Charge Description

3.13.3. Description

Self-contained summary of the *charge's* purpose and price.

3.13.4. Content Constraints

Constraint	Value
Column type	Dimension
Feature level	Mandatory
Allows nulls	True
Data type	String
Value format	<not specified>

3.13.5. Introduced (version)

1.0-preview

3.14. Charge Frequency

Charge Frequency indicates how often a [charge](#) will occur. Along with the [charge period](#) related columns, the Charge Frequency is commonly used to understand recurrence periods (e.g., monthly, yearly), forecast upcoming *charges*, and differentiate between one-time and recurring fees for purchases.

The ChargeFrequency column adheres to the following requirements:

- ChargeFrequency is RECOMMENDED to be present in a [FOCUS dataset](#).
- ChargeFrequency MUST be of type String.
- ChargeFrequency MUST NOT be null.
- ChargeFrequency MUST be one of the allowed values.
- ChargeFrequency MUST NOT be "Usage-Based" when [ChargeCategory](#) is "Purchase".

3.14.1. Column ID

ChargeFrequency

3.14.2. Display Name

Charge Frequency

3.14.3. Description

Indicates how often a *charge* will occur.

3.14.4. Content Constraints

Constraint	Value
Column type	Dimension
Feature level	Recommended
Allows nulls	False
Data type	String
Value format	Allowed values

Allowed values:

Value	Description
One-Time	<i>Charges</i> that only happen once and will not repeat. One-time <i>charges</i> are typically recorded on the hour or day when the cost was incurred.
Recurring	<i>Charges</i> that repeat on a periodic cadence (e.g., weekly, monthly) regardless of whether the product or service was used. Recurring <i>charges</i> typically happen on the same day or point within every period. The charge date does not change based on how or when the <i>service</i> is used.
Usage-Based	<i>Charges</i> that repeat every time the <i>service</i> is used. Usage-based <i>charges</i> are typically recorded hourly or daily, based on the granularity of the cost data for the period when the <i>service</i> was used (referred to as <i>charge period</i>). Usage-based <i>charges</i> are not recorded when the <i>service</i> is not used.

3.14.5. Introduced (version)

1.0-preview

3.15. Charge Period End

Charge Period End represents the [exclusive end bound](#) of a [charge period](#). For example, a time period where [Charge Period Start](#) is '2024-01-01T00:00:00Z' and Charge Period End is '2024-01-02T00:00:00Z' includes [charges](#) for January 1 since Charge Period Start represents the [inclusive start bound](#), but does not include *charges* for January 2 since Charge Period End represents the *exclusive end bound*.

The ChargePeriodEnd column adheres to the following requirements:

- ChargePeriodEnd MUST be present in a [FOCUS dataset](#).
- ChargePeriodEnd MUST be of type Date/Time.
- ChargePeriodEnd MUST conform to [DateTimeFormat](#) requirements.
- ChargePeriodEnd MUST NOT be null.
- ChargePeriodEnd MUST be the *exclusive end bound* of the effective period of the *charge*.

3.15.1. Column ID

ChargePeriodEnd

3.15.2. Display Name

Charge Period End

3.15.3. Description

The *exclusive end bound* of a *charge period*.

3.15.4. Content constraints

Constraint	Value
Column type	Dimension
Feature level	Mandatory
Allows nulls	False
Data type	Date/Time
Value format	Date/Time Format

3.15.5. Introduced (version)

0.5

3.16. Charge Period Start

Charge Period Start represents the [inclusive start bound](#) of a [charge period](#). For example, a time period where Charge Period Start is '2024-01-01T00:00:00Z' and [Charge Period End](#) is '2024-01-02T00:00:00Z' includes [charges](#) for January 1 since Charge Period Start represents the *inclusive start bound*, but does not include *charges* for January 2 since Charge Period End represents the [exclusive end bound](#).

The ChargePeriodStart column adheres to the following requirements:

- ChargePeriodStart MUST be present in a [FOCUS dataset](#).
- ChargePeriodStart MUST be of type Date/Time.
- ChargePeriodStart MUST conform to [DateTimeFormat](#) requirements.
- ChargePeriodStart MUST NOT be null.
- ChargePeriodStart MUST be the *inclusive start bound* of the effective period of the *charge*.

3.16.1. Column ID

ChargePeriodStart

3.16.2. Display Name

Charge Period Start

3.16.3. Description

The *inclusive start bound* of a *charge period*.

3.16.4. Content constraints

Constraint	Value
Column type	Dimension
Feature level	Mandatory
Allows nulls	False
Data type	Date/Time
Value format	Date/Time Format

3.16.5. Introduced (version)

0.5

3.17. Commitment Discount Category

Commitment Discount Category indicates whether the [commitment discount](#) identified in the CommitmentDiscountId column is based on usage quantity or cost (aka "spend"). The CommitmentDiscountCategory column is only applicable to *commitment discounts* and not

[negotiated discounts](#).

The CommitmentDiscountCategory column adheres to the following requirements:

- CommitmentDiscountCategory MUST be present in a [FOCUS dataset](#) when the provider supports *commitment discounts*.
- CommitmentDiscountCategory MUST be of type String.
- CommitmentDiscountCategory nullability is defined as follows:
 - CommitmentDiscountCategory MUST be null when [CommitmentDiscountId](#) is null.
 - CommitmentDiscountCategory MUST NOT be null when CommitmentDiscountId is not null.
- CommitmentDiscountCategory MUST be one of the allowed values.

3.17.1. Column ID

CommitmentDiscountCategory

3.17.2. Display Name

Commitment Discount Category

3.17.3. Description

Indicates whether the *commitment discount* identified in the CommitmentDiscountId column is based on usage quantity or cost (aka "spend").

3.17.4. Content constraints

Constraint	Value
Column type	Dimension
Feature level	Conditional
Allows nulls	True
Data type	String
Value format	Allowed Values

Allowed values:

Value	Description
Spend	Commitment discounts that require a predetermined amount of spend.
Usage	Commitment discounts that require a predetermined amount of usage.

3.17.5. Introduced (version)

1.0-preview

3.18. Commitment Discount ID

A Commitment Discount ID is the identifier assigned to a [commitment discount](#) by the provider. Commitment Discount ID is commonly used for scenarios like chargeback for *commitments* and savings per *commitment discount*. The CommitmentDiscountId column is only applicable to *commitment discounts* and not [negotiated discounts](#).

The CommitmentDiscountId column adheres to the following requirements:

- CommitmentDiscountId MUST be present in a [FOCUS dataset](#) when the provider supports *commitment discounts*.
- CommitmentDiscountId MUST be of type String.
- CommitmentDiscountId MUST conform to [StringHandling](#) requirements.
- CommitmentDiscountId nullability is defined as follows:
 - CommitmentDiscountId MUST be null when a [charge](#) is not related to a *commitment discount*.
 - CommitmentDiscountId MUST NOT be null when a *charge* is related to a *commitment discount*.
- When CommitmentDiscountId is not null, CommitmentDiscountId adheres to the following additional requirements:
 - CommitmentDiscountId MUST be a unique identifier within the provider.
 - CommitmentDiscountId SHOULD be a fully-qualified identifier.

3.18.1. Column ID

CommitmentDiscountId

3.18.2. Display Name

Commitment Discount ID

3.18.3. Description

The identifier assigned to a *commitment discount* by the provider.

3.18.4. Content constraints

Constraint	Value
Column type	Dimension
Feature level	Conditional
Allows nulls	True
Data type	String
Value format	<not specified>

3.18.5. Introduced (version)

1.0-preview

3.19. Commitment Discount Name

A Commitment Discount Name is the display name assigned to a [commitment discount](#). The CommitmentDiscountName column is only applicable to *commitment discounts* and not [negotiated discounts](#).

The CommitmentDiscountName column adheres to the following requirements:

- CommitmentDiscountName MUST be present in a [FOCUS dataset](#) when the provider supports *commitment discounts*.
- CommitmentDiscountName MUST be of type String.
- CommitmentDiscountName MUST conform to [StringHandling](#) requirements.
- CommitmentDiscountName nullability is defined as follows:
 - CommitmentDiscountName MUST be null when [CommitmentDiscountId](#) is null.
 - When CommitmentDiscountId is not null, CommitmentDiscountName adheres to the following additional requirements:
 - CommitmentDiscountName MUST NOT be null when a display name can be assigned to a *commitment discount*.
 - CommitmentDiscountName MAY be null when a display name cannot be assigned to a *commitment discount*.

3.19.1. Column ID

CommitmentDiscountName

3.19.2. Display Name

Commitment Discount Name

3.19.3. Description

The display name assigned to a *commitment discount*.

3.19.4. Content constraints

Constraint	Value
Column type	Dimension
Feature level	Conditional
Allows nulls	True
Data type	String
Value format	<not specified>

3.19.5. Introduced (version)

1.0-preview

3.20. Commitment Discount Quantity

Commitment Discount Quantity is the amount of a [commitment discount](#) purchased or accounted for in *commitment discount* related [rows](#) that is denominated in [Commitment Discount Units](#). The aggregated Commitment Discount Quantity across purchase records, pertaining to a particular [Commitment Discount ID](#) during its [term](#), represents the total Commitment Discount Units acquired with that commitment discount. For committed usage, the Commitment Discount Quantity is either the number of Commitment Discount Units consumed by a *row* that is covered by a *commitment discount* or is the unused portion of a *commitment discount* over a [charge period](#). Commitment Discount Quantity is commonly used in *commitment discount* analysis and optimization use cases and only applies to *commitment discounts*, not [negotiated discounts](#).

When [CommitmentDiscountCategory](#) is "Usage" (usage-based *commitment discounts*), the Commitment Discount Quantity reflects the predefined amount of usage purchased or consumed. If [commitment discount flexibility](#) is applicable, this value may be further transformed based on additional, provider-specific requirements. When [CommitmentDiscountCategory](#) is "Spend" (spend-based *commitment discounts*), the Commitment Discount Quantity reflects the predefined amount of spend purchased or consumed. See [Appendix: Commitment Discount Flexibility](#) for more details around *commitment discount flexibility*.

The CommitmentDiscountQuantity column adheres to the following requirements:

- CommitmentDiscountQuantity MUST be present in a [FOCUS dataset](#) when the provider supports *commitment discounts*.
- CommitmentDiscountQuantity MUST be of type Decimal.
- CommitmentDiscountQuantity MUST conform to [NumericFormat](#) requirements.
- CommitmentDiscountQuantity nullability is defined as follows:
 - When ChargeCategory is "Usage" or "Purchase" and CommitmentDiscountId is not null, CommitmentDiscountQuantity adheres to the following additional requirements:
 - CommitmentDiscountQuantity MUST NOT be null when [ChargeClass](#) is not "Correction".
 - CommitmentDiscountQuantity MAY be null when ChargeClass is "Correction".
 - CommitmentDiscountQuantity MUST be null in all other cases.
- When CommitmentDiscountQuantity is not null, CommitmentDiscountQuantity adheres to the following additional requirements:
 - CommitmentDiscountQuantity MUST be a valid decimal value.
 - When ChargeCategory is "Purchase":
 - CommitmentDiscountQuantity MUST be the quantity of CommitmentDiscountUnit, paid fully or partially upfront, that is eligible for consumption over the *commitment discount's term* when [ChargeFrequency](#) is "One-Time".
 - CommitmentDiscountQuantity MUST be the quantity of CommitmentDiscountUnit that is eligible for consumption for each *charge period* that corresponds with the purchase when ChargeFrequency is "Recurring".
 - When ChargeCategory is "Usage":
 - CommitmentDiscountQuantity MUST be the metered quantity of CommitmentDiscountUnit that is consumed in a given *charge period* when [CommitmentDiscountStatus](#) is "Used".
 - CommitmentDiscountQuantity MUST be the remaining, unused quantity of CommitmentDiscountUnit in a given *charge period* when CommitmentDiscountStatus is "Unused".

3.20.1. Column ID

CommitmentDiscountQuantity

3.20.2. Display Name

Commitment Discount Quantity

3.20.3. Description

The amount of a *commitment discount* purchased or accounted for in *commitment discount* related *rows* that is denominated in Commitment Discount Units.

3.20.4. Usability Constraints

Aggregation: When aggregating Commitment Discount Quantity for commitment utilization calculations, it's important to exclude [commitment discount](#) purchases (i.e. when Charge Category is "Purchase") that are paid to cover future eligible [charges](#) (e.g., *commitment discount*). Otherwise, when accounting for all upfront or accrued purchases, it's important to exclude *commitment discount* usage (i.e. when Charge Category is "Usage"). This exclusion helps prevent double counting of these quantities in the aggregation.

3.20.5. Content constraints

Constraint	Value
Column type	Metric
Feature level	Conditional
Allows nulls	True
Data type	Decimal
Value format	Numeric Format
Number range	Any valid decimal value

3.20.6. Introduced (version)

1.1

3.21. Commitment Discount Status

Commitment Discount Status indicates whether the [charge](#) corresponds with the consumption of a [commitment discount](#) identified in the CommitmentDiscountId column or the unused portion of the committed amount. The CommitmentDiscountStatus column is only applicable to *commitment discounts* and not [negotiated discounts](#).

The CommitmentDiscountStatus column adheres to the following requirements:

- CommitmentDiscountStatus MUST be present in a [FOCUS dataset](#) when the provider supports *commitment discounts*.
- CommitmentDiscountStatus MUST be of type String.
- CommitmentDiscountStatus nullability is defined as follows:
 - CommitmentDiscountStatus MUST be null when [CommitmentDiscountId](#) is null.
 - CommitmentDiscountStatus MUST NOT be null when CommitmentDiscountId is not null and [Charge Category](#) is "Usage".
- CommitmentDiscountStatus MUST be one of the allowed values.

3.21.1. Column ID

CommitmentDiscountStatus

3.21.2. Display name

Commitment Discount Status

3.21.3. Description

Indicates whether the *charge* corresponds with the consumption of a *commitment discount* or the unused portion of the committed amount.

3.21.4. Content constraints

Constraint	Value
Column type	Dimension
Feature level	Conditional
Allows nulls	True
Data type	String
Value format	Allowed Values

Allowed values:

Value	Description
Used	<i>Charges</i> that utilized a specific amount of a commitment discount.
Unused	<i>Charges</i> that represent the unused portion of the commitment discount.

3.21.5. Introduced (version)

1.0

3.22. Commitment Discount Type

Commitment Discount Type is a provider-assigned name to identify the type of [commitment discount](#) applied to the [row](#). The CommitmentDiscountType column is only applicable to *commitment discounts* and not [negotiated discounts](#).

The CommitmentDiscountType column adheres to the following requirements:

- CommitmentDiscountType MUST be present in a [FOCUS dataset](#) when the provider supports *commitment discounts*.
- CommitmentDiscountType MUST be of type String.
- CommitmentDiscountType MUST conform to [StringHandling](#) requirements.
- CommitmentDiscountType nullability is defined as follows:
 - CommitmentDiscountType MUST be null when [CommitmentDiscountId](#) is null.
 - CommitmentDiscountType MUST NOT be null when CommitmentDiscountId is not null.

3.22.1. Column ID

CommitmentDiscountType

3.22.2. Display Name

Commitment Discount Type

3.22.3. Description

A provider-assigned identifier for the type of *commitment discount* applied to the *row*.

3.22.4. Content Constraints

Constraint	Value
Column type	Dimension
Feature level	Conditional
Allows nulls	True
Data type	String
Value format	<not specified>

3.22.5. Introduced (version)

1.0-preview

3.23. Commitment Discount Unit

Commitment Discount Unit represents the provider-specified measurement unit indicating how a provider measures the [Commitment Discount Quantity](#) of a [commitment discount](#). The CommitmentDiscountUnit column is only applicable to *commitment discounts* and not [negotiated discounts](#).

The CommitmentDiscountUnit column adheres to the following requirements:

- CommitmentDiscountUnit MUST be present in a [FOCUS dataset](#) when the provider supports *commitment discounts*.
- CommitmentDiscountUnit MUST be of type String.
- CommitmentDiscountUnit MUST conform to [StringHandling](#) requirements.
- CommitmentDiscountUnit SHOULD conform to [UnitFormat](#) requirements.
- CommitmentDiscountUnit nullability is defined as follows:
 - CommitmentDiscountUnit MUST be null when CommitmentDiscountQuantity is null.
 - CommitmentDiscountUnit MUST NOT be null when CommitmentDiscountQuantity is not null.
- When CommitmentDiscountUnit is not null, CommitmentDiscountUnit adheres to the following additional requirements:
 - CommitmentDiscountUnit MUST remain consistent over time for a given

CommitmentDiscountId.

- CommitmentDiscountUnit MUST represent the unit used to measure the *commitment discount*.
- When accounting for [commitment discount flexibility](#), the CommitmentDiscountUnit value SHOULD reflect this consideration.

See [Examples: Commitment Discount Flexibility](#) for more details around *commitment discount flexibility*.

3.23.1. Column ID

CommitmentDiscountUnit

3.23.2. Display Name

Commitment Discount Unit

3.23.3. Description

The provider-specified measurement unit indicating how a provider measures the Commitment Discount Quantity of a *commitment discount*.

3.23.4. Content constraints

Constraint	Value
Column type	Dimension
Feature level	Conditional
Allows nulls	True
Data type	String
Value format	Unit Format

3.23.5. Introduced (version)

1.1

3.24. Consumed Quantity

The Consumed Quantity represents the volume of a metered SKU associated with a [resource](#) or [service](#) used, based on the [Consumed Unit](#). Consumed Quantity is often derived at a finer granularity or over a different time interval when compared to the [Pricing Quantity](#) (complementary to [Pricing Unit](#)) and focuses on *resource* and *service* consumption, not pricing and cost.

The ConsumedQuantity column adheres to the following requirements:

- ConsumedQuantity MUST be present in a [FOCUS dataset](#) when the provider supports the measurement of usage.

- ConsumedQuantity MUST be of type Decimal.
- ConsumedQuantity MUST conform to [NumericFormat](#) requirements.
- ConsumedQuantity nullability is defined as follows:
 - ConsumedQuantity MUST be null when [ChargeCategory](#) is not "Usage", or when ChargeCategory is "Usage" and [CommitmentDiscountStatus](#) is "Unused".
 - When ChargeCategory is "Usage" and CommitmentDiscountStatus is not "Unused", ConsumedQuantity adheres to the following additional requirements:
 - ConsumedQuantity MUST NOT be null when [ChargeClass](#) is not "Correction".
 - ConsumedQuantity MAY be null when ChargeClass is "Correction".
- ConsumedQuantity MUST be a valid decimal value when not null.

3.24.1. Column ID

ConsumedQuantity

3.24.2. Display Name

Consumed Quantity

3.24.3. Description

The volume of a metered SKU associated with a *resource* or *service* used, based on the Consumed Unit.

3.24.4. Content constraints

Constraint	Value
Column type	Metric
Feature level	Conditional
Allows nulls	True
Data type	Decimal
Value format	Numeric Format
Number range	Any valid decimal value

3.24.5. Introduced (version)

1.0

3.25. Consumed Unit

The Consumed Unit represents a provider-specified measurement unit indicating how a provider measures usage of a metered SKU associated with a [resource](#) or [service](#). Consumed Unit complements the [Consumed Quantity](#) metric. It is often listed at a finer granularity or over a different time interval when compared to [Pricing Unit](#) (complementary to [Pricing Quantity](#)), and focuses on *resource* and *service* consumption, not pricing and cost.

The ConsumedUnit column adheres to the following requirements:

- ConsumedUnit MUST be present in a [FOCUS dataset](#) when the provider supports the measurement of usage.
- ConsumedUnit MUST be of type String.
- ConsumedUnit MUST conform to [StringHandling](#) requirements.
- ConsumedUnit SHOULD conform to [UnitFormat](#) requirements.
- ConsumedUnit nullability is defined as follows:
 - ConsumedUnit MUST be null when ConsumedQuantity is null.
 - ConsumedUnit MUST NOT be null when ConsumedQuantity is not null.

3.25.1. Column ID

ConsumedUnit

3.25.2. Display Name

Consumed Unit

3.25.3. Description

Provider-specified measurement unit indicating how a provider measures usage of a metered SKU associated with a *resource* or *service*.

3.25.4. Content constraints

Constraint	Value
Column type	Dimension
Feature level	Conditional
Allows nulls	True
Data type	String
Value format	Unit Format recommended

3.25.5. Introduced (version)

1.0

3.26. Contracted Cost

Contracted Cost represents the cost calculated by multiplying [contracted unit price](#) and the corresponding [Pricing Quantity](#). Contracted Cost is denominated in the [Billing Currency](#) and is commonly used for calculating savings based on negotiation activities, by comparing it with [List Cost](#). If [negotiated discounts](#) are not applicable, the Contracted Cost defaults to the List Cost.

The ContractedCost column adheres to the following requirements:

- ContractedCost MUST be present in a [FOCUS dataset](#).
- ContractedCost MUST be of type Decimal.
- ContractedCost MUST conform to [NumericFormat](#) requirements.
- ContractedCost MUST NOT be null.
- ContractedCost MUST be a valid decimal value.
- ContractedCost MUST be denominated in the BillingCurrency.
- When [ContractedUnitPrice](#) is null, ContractedCost adheres to the following additional requirements:
 - ContractedCost of a [charge](#) calculated based on other *charges* (e.g., when the [ChargeCategory](#) is "Tax") MUST be calculated based on the ContractedCost of those related *charges*.
 - ContractedCost of a *charge* unrelated to other *charges* (e.g., when the ChargeCategory is "Credit") MUST match the [BilledCost](#).
- The product of ContractedUnitPrice and PricingQuantity MUST match the ContractedCost when ContractedUnitPrice is not null, PricingQuantity is not null, and [ChargeClass](#) is not "Correction".
- Discrepancies in ContractedCost, ContractedUnitPrice, or PricingQuantity MAY exist when ChargeClass is "Correction".

3.26.1. Column ID

ContractedCost

3.26.2. Display Name

Contracted Cost

3.26.3. Description

Cost calculated by multiplying *contracted unit price* and the corresponding Pricing Quantity.

3.26.4. Usability Constraints

Aggregation: When aggregating Contracted Cost for savings calculations, it's important to exclude either [Charge Category](#) "Purchase" *charges* (one-time and recurring) that are paid to cover future eligible *charges* (e.g., [commitment discount](#)) or the covered [Charge Category](#) "Usage" *charges* themselves. This exclusion helps prevent double counting of these *charges* in the aggregation. Which set of *charges* to exclude depends on whether cost are aggregated on a billed basis (exclude covered *charges*) or accrual basis (exclude Purchases for future *charges*). For instance, *charges* categorized as [Charge Category](#) "Purchase" and their related [Charge Category](#) "Tax" *charges* for a Commitment Discount might be excluded from an accrual basis cost aggregation of Contracted Cost. This is because the "Usage" and "Tax" charge records provided during the term of the commitment discount already specify the Contracted Cost. Purchase *charges* that cover future eligible *charges* can be identified by filtering for [Charge Category](#) "Purchase" records with a [Billed Cost](#) greater than 0 and an [Effective Cost](#) equal to 0.

3.26.5. Content Constraints

Constraint	Value
------------	-------

Constraint	Value
Column type	Metric
Feature level	Mandatory
Allows nulls	False
Data type	Decimal
Value format	Numeric Format
Number range	Any valid decimal value

3.26.6. Introduced (version)

1.0

3.27. Contracted Unit Price

The Contracted Unit Price represents the agreed-upon unit price for a single [Pricing Unit](#) of the associated SKU, inclusive of [negotiated discounts](#), if present, while excluding negotiated [commitment discounts](#) or any other discounts. This price is denominated in the [Billing Currency](#). The Contracted Unit Price is commonly used for calculating savings based on negotiation activities. If negotiated discounts are not applicable, the Contracted Unit Price defaults to the [List Unit Price](#).

The ContractedUnitPrice column adheres to the following requirements:

- ContractedUnitPrice MUST be present in a [FOCUS dataset](#) when the provider supports negotiated pricing concepts.
- ContractedUnitPrice adheres to the following additional requirements:
- ContractedUnitPrice MUST be of type Decimal.
- ContractedUnitPrice MUST conform to [NumericFormat](#) requirements.
- ContractedUnitPrice nullability is defined as follows:
 - ContractedUnitPrice MUST be null when [ChargeCategory](#) is "Tax".
 - ContractedUnitPrice MUST NOT be null when ChargeCategory is "Usage" or "Purchase" and [ChargeClass](#) is not "Correction".
 - ContractedUnitPrice MAY be null in all other cases.
- When ContractedUnitPrice is not null, ContractedUnitPrice adheres to the following additional requirements:
 - ContractedUnitPrice MUST be a non-negative decimal value.
 - ContractedUnitPrice MUST be denominated in the BillingCurrency.
 - The product of ContractedUnitPrice and [PricingQuantity](#) MUST match the [ContractedCost](#) when PricingQuantity is not null and ChargeClass is not "Correction".
- Discrepancies in ContractedUnitPrice, ContractedCost, or PricingQuantity MAY exist when ChargeClass is "Correction".

3.27.1. Column ID

ContractedUnitPrice

3.27.2. Display Name

Contracted Unit Price

3.27.3. Description

The agreed-upon unit price for a single Pricing Unit of the associated SKU, inclusive of negotiated discounts, if present, while excluding negotiated commitment discounts or any other discounts.

3.27.4. Usability Constraints

Aggregation: Column values should only be viewed in the context of their row and not aggregated to produce a total.

3.27.5. Content Constraints

Constraint	Value
Column type	Metric
Feature level	Conditional
Allows nulls	True
Data type	Decimal
Value format	Numeric Format
Number range	Any valid non-negative decimal value

3.27.6. Introduced (version)

1.0

3.28. Effective Cost

Effective Cost represents the [amortized](#) cost of the [charge](#) after applying all reduced rates, discounts, and the applicable portion of relevant, prepaid purchases (one-time or recurring) that covered this [charge](#). The [amortized](#) portion included should be proportional to the [Pricing Quantity](#) and the time granularity of the data. Since amortization breaks down and spreads the cost of a prepaid purchase, to subsequent eligible [charges](#), the Effective Cost of the original prepaid [charge](#) is set to 0. Effective Cost does not mix or "blend" costs across multiple [charges](#) of the same [service](#). This cost is denominated in the [Billing Currency](#). The Effective Cost is commonly utilized to track and analyze spending trends.

This column resolves two challenges that are faced by practitioners:

1. Practitioners need to [amortize](#) relevant purchases, such as upfront fees, throughout the [commitment](#) and distribute them to the appropriate reporting groups (e.g., [tags](#), [resources](#)).
2. Many [commitment discount](#) constructs include a recurring expense for the [commitment](#) for every [billing period](#) and must distribute this cost to the [resources](#) using the [commitment](#). This forces reconciliation between the initial [commitment row](#) per period and the actual usage [rows](#).

The EffectiveCost column adheres to the following requirements:

- EffectiveCost MUST be present in a [FOCUS dataset](#).
- EffectiveCost MUST be of type Decimal.
- EffectiveCost MUST conform to [NumericFormat](#) requirements.
- EffectiveCost MUST NOT be null.
- EffectiveCost MUST be a valid decimal value.
- EffectiveCost MUST be 0 when [ChargeCategory](#) is "Purchase" and the purchase is intended to

- cover future eligible *charges*.
- EffectiveCost MUST be denominated in the BillingCurrency.
- The sum of EffectiveCost in a given *billing period* may not match the sum of the invoices received for the same *billing period* for a [billing account](#).
- When ChargeCategory is not "Usage" or "Purchase", EffectiveCost adheres to the following additional requirements:
 - EffectiveCost of a *charge* calculated based on other *charges* (e.g., when the ChargeCategory is "Tax") MUST be calculated based on the EffectiveCost of those related *charges*.
 - EffectiveCost of a *charge* unrelated to other *charges* (e.g., when the ChargeCategory is "Credit") MUST match the [BilledCost](#).
- *Charges* for a given [CommitmentDiscountId](#) adhere to the following additional requirements:
 - The sum of EffectiveCost where ChargeCategory is "Usage" MUST equal the sum of BilledCost where ChargeCategory is "Purchase".
 - The sum of EffectiveCost where ChargeCategory is "Usage" MUST equal the sum of EffectiveCost where ChargeCategory is "Usage" and [CommitmentDiscountStatus](#) is "Used", plus the sum of EffectiveCost where ChargeCategory is "Usage" and CommitmentDiscountStatus is "Unused".

3.28.1. Column ID

EffectiveCost

3.28.2. Display Name

Effective Cost

3.28.3. Description

The *amortized* cost of the *charge* after applying all reduced rates, discounts, and the applicable portion of relevant, prepaid purchases (one-time or recurring) that covered this *charge*.

3.28.3.1. Concerning Granularity and Distribution of Recurring Fee

Providers should distribute the *commitment* purchase amount instead of including a *row* at the beginning of a period so practitioners do not need to manually distribute the fee themselves.

3.28.3.2. Concerning Amortization Approaches

Eligible purchases should be *amortized* using a methodology determined by the provider that reflects the needs of their customer base and is proportional to the Pricing Quantity and the time granularity of the *row*. Should a practitioner desire to *amortize* relevant purchases using a different approach, the practitioner can do so using the [Billed Cost](#) for the line item representing the initial purchase.

3.28.4. Content constraints

Constraint	Value
Column type	Metric
Feature level	Mandatory
Allows nulls	False
Data type	Decimal
Value format	Numeric Format
Number range	Any valid decimal value

3.28.5. Introduced (version)

0.5

3.29. Invoice ID

An Invoice ID is a provider-assigned identifier for an invoice encapsulating some or all [charges](#) in the corresponding [billing period](#) for a given [billing account](#). Invoices are commonly used for scenarios like tracking billing transactions, facilitating payment processes and for performing invoice reconciliation between *charges* and billing periods.

The InvoiceId column adheres to the following requirements:

- InvoiceId is RECOMMENDED to be present in a [FOCUS dataset](#).
- InvoiceId MUST be of type String.
- InvoiceId MUST conform to [StringHandling](#) requirements.
- InvoiceId nullability is defined as follows:
 - InvoiceId MUST be null when the *charge* is not associated either with an invoice or with a pre-generated provisional invoice.
 - InvoiceId MUST NOT be null when the *charge* is associated with either an issued invoice or a pre-generated provisional invoice.
- InvoiceId MAY be generated prior to an invoice being issued.
- InvoiceId MUST be associated with the related *charge* and BillingAccountId when a pre-generated invoice or provisional invoice exists.

See [Appendix: Grouping constructs for resources or services](#) for details and examples of the different grouping constructs supported by FOCUS.

3.29.1. Column ID

InvoiceId

3.29.2. Display Name

Invoice ID

3.29.3. Description

The provider-assigned identifier for an invoice encapsulating some or all *charges* in the corresponding billing period for a given billing account.

3.29.4. Content constraints

Constraint	Value
Column type	Dimension
Feature level	Recommended
Allows nulls	True
Data type	String
Value format	<not specified>

3.29.5. Introduced (version)

1.2

3.30. Invoice Issuer

An Invoice Issuer is an entity responsible for issuing payable invoices for the [resources](#) or [services](#) consumed. It is commonly used for cost analysis and reporting scenarios.

The InvoiceIssuerName column adheres to the following requirements:

- InvoiceIssuerName MUST be present in a [FOCUS dataset](#).
- InvoiceIssuerName MUST be of type String.
- InvoiceIssuerName MUST conform to [StringHandling](#) requirements.
- InvoiceIssuerName MUST NOT be null.

See [Appendix: Origination of cost data](#) section for examples of [Provider](#), [Publisher](#) and Invoice Issuer values that can be used for various purchasing scenarios.

3.30.1. Column ID

InvoiceIssuerName

3.30.2. Display Name

Invoice Issuer

3.30.3. Description

The name of the entity responsible for invoicing for the *resources* or *services* consumed.

3.30.4. Content Constraints

Constraint	Value
Column type	Dimension

Constraint	Value
Feature level	Mandatory
Allows nulls	False
Data type	String
Value format	<not specified>

3.30.5. Introduced (version)

0.5

3.31. List Cost

List Cost represents the cost calculated by multiplying the [list unit price](#) and the corresponding [Pricing Quantity](#). List Cost is denominated in the [Billing Currency](#) and is commonly used for calculating savings based on various rate optimization activities by comparing it with [Contracted Cost](#), [Billed Cost](#) and [Effective Cost](#).

The ListCost column adheres to the following requirements:

- ListCost MUST be present in a [FOCUS dataset](#).
- ListCost MUST be of type Decimal.
- ListCost MUST conform to [NumericFormat](#) requirements.
- ListCost MUST NOT be null.
- ListCost MUST be a valid decimal value.
- ListCost MUST be denominated in the BillingCurrency.
- When [ListUnitPrice](#) is null, ListCost adheres to the following additional requirements:
 - ListCost of a [charge](#) calculated based on other *charges* (e.g., when the [ChargeCategory](#) is "Tax") MUST be calculated based on the ListCost of those related *charges*.
 - ListCost of a *charge* unrelated to other *charges* (e.g., when the ChargeCategory is "Credit") MUST match the [BilledCost](#).
- The product of ListUnitPrice and PricingQuantity MUST match the ListCost when ListUnitPrice is not null, PricingQuantity is not null, and [ChargeClass](#) is not "Correction".
- Discrepancies in ListCost, ListUnitPrice, or PricingQuantity MAY exist when ChargeClass is "Correction".

3.31.1. Column ID

ListCost

3.31.2. Display Name

List Cost

3.31.3. Description

Cost calculated by multiplying List Unit Price and the corresponding Pricing Quantity.

3.31.4. Usability Constraints

Aggregation: When aggregating List Cost for savings calculations, it's important to exclude either [Charge Category](#) "Purchase" *charges* (one-time and recurring) that are paid to cover future eligible *charges* (e.g., [commitment discount](#)) or the covered [Charge Category](#) "Usage" *charges* themselves. This exclusion helps prevent double counting of these *charges* in the aggregation. Which set of *charges* to exclude depends on whether cost are aggregated on a billed basis (exclude covered *charges*) or accrual basis (exclude Purchases for future *charges*). For instance, *charges* categorized as [Charge Category](#) "Purchase" and their related [Charge Category](#) "Tax" *charges* for a Commitment Discount might be excluded from an accrual basis cost aggregation of List Cost. This is because the "Usage" and "Tax" charge records provided during the term of the commitment discount already specify the List Cost. Purchase *charges* that cover future eligible *charges* can be identified by filtering for [Charge Category](#) "Purchase" records with a [Billed Cost](#) greater than 0 and an [Effective Cost](#) equal to 0.

3.31.5. Content Constraints

Constraint	Value
Column type	Metric
Feature level	Mandatory
Allows nulls	False
Data type	Decimal
Value format	Numeric Format
Number range	Any valid decimal value

3.31.6. Introduced (version)

1.0-preview

3.32. List Unit Price

The List Unit Price represents the suggested provider-published unit price for a single [Pricing Unit](#) of the associated SKU, exclusive of any discounts. This price is denominated in the [Billing Currency](#). The List Unit Price is commonly used for calculating savings based on various rate optimization activities.

The ListUnitPrice column adheres to the following requirements:

- ListUnitPrice MUST be present in a [FOCUS dataset](#) when the provider publishes unit prices exclusive of discounts.
- ListUnitPrice MUST be of type Decimal.
- ListUnitPrice MUST conform to [NumericFormat](#) requirements.
- ListUnitPrice nullability is defined as follows:
 - ListUnitPrice MUST be null when [ChargeCategory](#) is "Tax".
 - ListUnitPrice MUST NOT be null when ChargeCategory is "Usage" or "Purchase" and [ChargeClass](#) is not "Correction".
 - ListUnitPrice MAY be null in all other cases.
- When ListUnitPrice is not null, ListUnitPrice adheres to the following additional requirements:
 - ListUnitPrice MUST be a non-negative decimal value.
 - ListUnitPrice MUST be denominated in the BillingCurrency.
 - The product of ListUnitPrice and [PricingQuantity](#) MUST match the [ListCost](#) when PricingQuantity is not null and ChargeClass is not "Correction".
 - Discrepancies in ListUnitPrice, ListCost, or PricingQuantity MAY exist when ChargeClass is

"Correction".

3.32.1. Column ID

ListUnitPrice

3.32.2. Display Name

List Unit Price

3.32.3. Description

The suggested provider-published unit price for a single Pricing Unit of the associated SKU, exclusive of any discounts.

3.32.4. Usability Constraints

Aggregation: Column values should only be viewed in the context of their row and not aggregated to produce a total.

3.32.5. Content Constraints

Constraint	Value
Column type	Metric
Feature level	Conditional
Allows nulls	True
Data type	Decimal
Value format	Numeric Format
Number range	Any valid non-negative decimal value

3.32.6. Introduced (version)

1.0-preview

3.33. Pricing Category

Pricing Category describes the pricing model used for a [charge](#) at the time of use or purchase. It can be useful for distinguishing between *charges* incurred at the [list unit price](#) or a reduced price and exposing optimization opportunities, like increasing [commitment discount](#) coverage.

The PricingCategory column adheres to the following requirements:

- PricingCategory MUST be present in a [FOCUS dataset](#) when the provider supports more than one pricing category across all [SKUs](#).

- PricingCategory MUST be of type String.
- PricingCategory nullability is defined as follows:
 - PricingCategory MUST be null when [ChargeCategory](#) is "Tax".
 - PricingCategory MUST NOT be null when ChargeCategory is "Usage" or "Purchase" and [ChargeClass](#) is not "Correction".
 - PricingCategory MAY be null in all other cases.
- When PricingCategory is not null, PricingCategory adheres to the following additional requirements:
 - PricingCategory MUST be one of the allowed values.
 - PricingCategory MUST be "Standard" when pricing is predetermined at the agreed upon rate for the [billing account](#).
 - PricingCategory MUST be "Committed" when the *charge* is subject to an existing *commitment discount* and is not the purchase of the *commitment discount*.
 - PricingCategory MUST be "Dynamic" when pricing is determined by the provider and may change over time, regardless of predetermined agreement pricing.
 - PricingCategory MUST be "Other" when there is a pricing model but none of the allowed values apply.

3.33.1. Column ID

PricingCategory

3.33.2. Display Name

Pricing Category

3.33.3. Description

Describes the pricing model used for a *charge* at the time of use or purchase.

3.33.4. Content constraints

Constraint	Value
Column type	Dimension
Feature level	Conditional
Allows nulls	True
Data type	String
Value format	Allowed values

Allowed values:

Value	Description
Standard	<i>Charges</i> priced at the agreed upon rate for the billing account, including negotiated discounts . This pricing includes any flat rate and volume/tiered pricing but does not include dynamic pricing or reduced pricing due to the application of a <i>commitment discount</i> . This does include the purchase of a commitment discount at agreed upon rates.

Value	Description
Dynamic	<i>Charges</i> priced at a variable rate determined by the provider. This includes any product or service with a unit price the provider can change without notice, like interruptible or low priority resources .
Committed	<i>Charges</i> with reduced pricing due to the application of the <i>commitment discount</i> specified by the Commitment Discount ID.
Other	<i>Charges</i> priced in a way not covered by another pricing category.

3.33.5. Introduced (version)

1.0-preview

3.34. Pricing Currency

[Pricing Currency](#) is the national or virtual currency denomination that a [resource](#) or [service](#) was priced in. Pricing Currency is commonly used in scenarios where different currencies are used for pricing and billing.

The PricingCurrency column adheres to the following requirements:

- PricingCurrency MUST be present in a [FOCUS dataset](#) when the provider supports pricing and billing in different currencies.
- PricingCurrency MUST be of type String.
- PricingCurrency MUST conform to [StringHandling](#) requirements.
- PricingCurrency MUST conform to [CurrencyFormat](#) requirements.
- PricingCurrency MUST NOT be null.

3.34.1. Column ID

PricingCurrency

3.34.2. Display Name

Pricing Currency

3.34.3. Description

The national or virtual currency denomination that a *resource* or *service* was priced in.

3.34.4. Content Constraints

Constraint	Value
Column type	Dimension
Feature level	Conditional
Allows nulls	True
Data type	String

Constraint	Value
Value format	Currency Format

3.34.5. Introduced (version)

1.2

3.35. Pricing Currency Contracted Unit Price

The Pricing Currency Contracted Unit Price represents the agreed-upon unit price for a single [Pricing Unit](#) of the associated [SKU](#), inclusive of [negotiated discounts](#), if present, while excluding negotiated [commitment discounts](#) or any other discounts. This price is denominated in the [Pricing Currency](#). When negotiated discounts do not apply to unit prices and instead are applied to exchange rates, the Pricing Currency Contracted Unit Price defaults to the [Pricing Currency List Unit Price](#). The Pricing Currency Contracted Unit Price is commonly used to calculate savings based on negotiation activities.

The PricingCurrencyContractedUnitPrice column adheres to the following requirements:

- PricingCurrencyContractedUnitPrice presence in a [FOCUS dataset](#) is defined as follows:
 - PricingCurrencyContractedUnitPrice MUST be present in a *FOCUS dataset* when the provider supports prices in virtual currency and publishes unit prices exclusive of discounts.
 - PricingCurrencyContractedUnitPrice is RECOMMENDED to be present in a *FOCUS dataset* when the provider supports pricing and billing in different currencies and publishes unit prices exclusive of discounts.
 - PricingCurrencyContractedUnitPrice MAY be present in a *FOCUS dataset* in all other cases.
- PricingCurrencyContractedUnitPrice MUST be of type Decimal.
- PricingCurrencyContractedUnitPrice MUST conform to [NumericFormat](#) requirements.
- PricingCurrencyContractedUnitPrice nullability is defined as follows:
 - PricingCurrencyContractedUnitPrice MUST be null when [ChargeCategory](#) is "Tax".
 - PricingCurrencyContractedUnitPrice MUST NOT be null when ChargeCategory is "Usage" or "Purchase" and [ChargeClass](#) is not "Correction".
 - PricingCurrencyContractedUnitPrice MAY be null in all other cases.
- When PricingCurrencyContractedUnitPrice is not null, PricingCurrencyContractedUnitPrice adheres to the following additional requirements:
 - PricingCurrencyContractedUnitPrice MUST be a non-negative decimal value.
 - PricingCurrencyContractedUnitPrice MUST be denominated in the PricingCurrency.
- Discrepancies in PricingCurrencyContractedUnitPrice, [ContractedCost](#), or [PricingQuantity](#) MAY exist when ChargeClass is "Correction".

3.35.1. Column ID

PricingCurrencyContractedUnitPrice

3.35.2. Display Name

Pricing Currency Contracted Unit Price

3.35.3. Description

The agreed-upon unit price for a single Pricing Unit of the associated SKU, inclusive of *negotiated discounts*, if present, while excluding negotiated *commitment discounts* or any other discounts, and expressed in Pricing Currency.

3.35.4. Usability Constraints

Aggregation: Column values should only be viewed in the context of their row and not aggregated to produce a total.

3.35.5. Content Constraints

Constraint	Value
Column type	Metric
Feature level	Conditional
Allows nulls	True
Data type	Decimal
Value format	Numeric Format
Number range	Any valid non-negative decimal value

3.35.6. Introduced (version)

1.2

3.36. Pricing Currency Effective Cost

The Pricing Currency Effective Cost represents the cost of the [charge](#) after applying all reduced rates, discounts, and the applicable portion of relevant, prepaid purchases (one-time or recurring) that covered this *charge*, as denominated in [Pricing Currency](#). This allows the practitioner to perform a conversion from either 1) a [national currency](#) to a [virtual currency](#) (e.g., tokens to USD), or 2) one national currency to another (e.g., EUR to USD).

The PricingCurrencyEffectiveCost column adheres to the following requirements:

- PricingCurrencyEffectiveCost presence in a [FOCUS dataset](#) is defined as follows:
 - PricingCurrencyEffectiveCost MUST be present in a *FOCUS dataset* when the provider supports prices in virtual currency and publishes unit prices exclusive of discounts.
 - PricingCurrencyEffectiveCost is RECOMMENDED to be present in a *FOCUS dataset* when the provider supports pricing and billing in different currencies and publishes unit prices exclusive of discounts.
 - PricingCurrencyEffectiveCost MAY be present in a *FOCUS dataset* in all other cases.
- PricingCurrencyEffectiveCost MUST be of type Decimal.
- PricingCurrencyEffectiveCost MUST conform to [NumericFormat](#) requirements.
- PricingCurrencyEffectiveCost MUST NOT be null.
- PricingCurrencyEffectiveCost MUST be a valid decimal value.
- PricingCurrencyEffectiveCost MUST be 0 in the event of prepaid purchases or purchases that are applicable to previous usage.
- PricingCurrencyEffectiveCost MUST be denominated in the [PricingCurrency](#).

3.36.1. Column ID

PricingCurrencyEffectiveCost

3.36.2. Display Name

Pricing Currency Effective Cost

3.36.3. Description

The cost of the *charge* after applying all reduced rates, discounts, and the applicable portion of relevant, prepaid purchases (one-time or recurring) that covered this *charge*, as denominated in Pricing Currency.

3.36.4. Content Constraints

Constraint	Value
Column type	Metric
Feature level	Conditional
Allows nulls	True
Data type	Decimal
Value format	Numeric Format
Number range	Any valid decimal value

3.36.5. Introduced (version)

1.2

3.37. Pricing Currency List Unit Price

The Pricing Currency List Unit Price represents the suggested provider-published unit price for a single [Pricing Unit](#) of the associated [SKU](#), exclusive of any discounts. This price is denominated in the [Pricing Currency](#). The Pricing Currency List Unit Price is commonly used for calculating savings based on various rate optimization activities.

The PricingCurrencyListUnitPrice column adheres to the following requirements:

- PricingCurrencyListUnitPrice presence in a [FOCUS dataset](#) is defined as follows:
 - PricingCurrencyListUnitPrice MUST be present in a *FOCUS dataset* when the provider supports prices in virtual currency and publishes unit prices exclusive of discounts.
 - PricingCurrencyListUnitPrice is RECOMMENDED to be present in a *FOCUS dataset* when the provider supports pricing and billing in different currencies and publishes unit prices exclusive of discounts.
 - PricingCurrencyListUnitPrice MAY be present in a *FOCUS dataset* in all other cases.
- PricingCurrencyListUnitPrice MUST be of type Decimal.

- PricingCurrencyListUnitPrice MUST conform to [NumericFormat](#) requirements.
- PricingCurrencyListUnitPrice nullability is defined as follows:
 - PricingCurrencyListUnitPrice MUST be null when [ChargeCategory](#) is "Tax".
 - PricingCurrencyListUnitPrice MUST NOT be null when ChargeCategory is "Usage" or "Purchase" and [ChargeClass](#) is not "Correction".
 - PricingCurrencyListUnitPrice MAY be null in all other cases.
- When PricingCurrencyListUnitPrice is not null, ListUnitPrice adheres to the following additional requirements:
 - PricingCurrencyListUnitPrice MUST be a non-negative decimal value.
 - PricingCurrencyListUnitPrice MUST be denominated in the PricingCurrency.
 - Discrepancies in PricingCurrencyListUnitPrice, ListCost, or PricingQuantity MAY be addressed independently when ChargeClass is "Correction".

3.37.1. Column ID

PricingCurrencyListUnitPrice

3.37.2. Display Name

Pricing Currency List Unit Price

3.37.3. Description

The suggested provider-published unit price for a single Pricing Unit of the associated *SKU*, exclusive of any discounts and expressed in Pricing Currency.

3.37.4. Usability Constraints

Aggregation: Column values should only be viewed in the context of their row and not aggregated to produce a total.

3.37.5. Content Constraints

Constraint	Value
Column type	Metric
Feature level	Conditional
Allows nulls	True
Data type	Decimal
Value format	Numeric Format
Number range	Any valid non-negative decimal value

3.37.6. Introduced (version)

1.2

3.38. Pricing Quantity

The Pricing Quantity represents the volume of a given [SKU](#) associated with a [resource](#) or [service](#) used or purchased, based on the [Pricing Unit](#). Distinct from [Consumed Quantity](#) (complementary to [Consumed Unit](#)), it focuses on pricing and cost, not *resource* and *service* consumption.

The PricingQuantity column adheres to the following requirements:

- PricingQuantity MUST be present in a [FOCUS dataset](#).
- PricingQuantity MUST be of type Decimal.
- PricingQuantity MUST conform to [NumericFormat](#) requirements.
- PricingQuantity nullability is defined as follows:
 - PricingQuantity MUST be null when [ChargeCategory](#) is "Tax".
 - PricingQuantity MUST NOT be null when ChargeCategory is "Usage" or "Purchase" and [ChargeClass](#) is not "Correction".
 - PricingQuantity MAY be null in all other cases.
- When PricingQuantity is not null, PricingQuantity adheres to the following additional requirements:
 - PricingQuantity MUST be a valid decimal value.
 - The product of PricingQuantity and a unit price (e.g., [ContractedUnitPrice](#)) MUST match the corresponding cost metric (e.g., [ContractedCost](#)) when the unit price is not null and ChargeClass is not "Correction".
- Discrepancies in PricingQuantity, unit prices (e.g., [ContractedUnitPrice](#)), or costs (e.g., [ContractedCost](#)) MAY exist when ChargeClass is "Correction".

3.38.1. Column ID

PricingQuantity

3.38.2. Display Name

Pricing Quantity

3.38.3. Description

The volume of a given *SKU* associated with a *resource* or *service* used or purchased, based on the Pricing Unit.

3.38.4. Usability Constraints

Aggregation: When aggregating Pricing Quantity for commitment utilization calculations, it's important to exclude [commitment discount](#) purchases (i.e. when Charge Category is "Purchase") that are paid to cover future eligible [charges](#) (e.g., *commitment discount*). Otherwise, when accounting for all upfront or accrued purchases, it's important to exclude *commitment discount* usage (i.e. when Charge Category is "Usage"). This exclusion helps prevent double counting of these quantities in the aggregation.

3.38.5. Content Constraints

Constraint	Value
Column type	Metric
Feature level	Mandatory
Allows nulls	True
Data type	Decimal
Value format	Numeric Format
Number Range	Any valid decimal value

3.38.6. Introduced (version)

1.0-preview

3.39. Pricing Unit

The Pricing Unit represents a provider-specified measurement unit for determining unit prices, indicating how the provider rates measured usage and purchase quantities after applying pricing rules like [block pricing](#). Common examples include the number of hours for compute appliance runtime (e.g., Hours), gigabyte-hours for a storage appliance (e.g., GB-Hours), or an accumulated count of requests for a network appliance or API service (e.g., 1000 Requests). Pricing Unit complements the [Pricing Quantity](#) metric. Distinct from the [Consumed Unit](#), it focuses on pricing and cost, not [resource](#) and [service](#) consumption, often at a coarser granularity.

The PricingUnit column adheres to the following requirements:

- PricingUnit MUST be present in a [FOCUS dataset](#).
- PricingUnit MUST be of type String.
- PricingUnit MUST conform to [StringHandling](#) requirements.
- PricingUnit SHOULD conform to [UnitFormat](#) requirements.
- PricingUnit nullability is defined as follows:
 - PricingUnit MUST be null when PricingQuantity is null.
 - PricingUnit MUST NOT be null when PricingQuantity is not null.
- When PricingUnit is not null, PricingUnit adheres to the following additional requirements:
 - PricingUnit MUST be semantically equal to the corresponding pricing measurement unit provided in provider-published [price list](#).
 - PricingUnit MUST be semantically equal to the corresponding pricing measurement unit provided in invoice, when the invoice includes a pricing measurement unit.

3.39.1. Column ID

PricingUnit

3.39.2. Display Name

Pricing Unit

3.39.3. Description

Provider-specified measurement unit for determining unit prices, indicating how the provider rates

measured usage and purchase quantities after applying pricing rules like *block pricing*.

3.39.4. Content constraints

Constraint	Value
Column type	Dimension
Feature level	Mandatory
Allows nulls	True
Data type	String
Value format	Unit Format

3.39.5. Introduced (version)

1.0-preview

3.40. Provider

A Provider is an entity that makes the [resources](#) or [services](#) available for purchase. It is commonly used for cost analysis and reporting scenarios.

The ProviderName column adheres to the following requirements:

- ProviderName MUST be present in a [FOCUS dataset](#).
- ProviderName MUST be of type String.
- ProviderName MUST conform to [StringHandling](#) requirements.
- ProviderName MUST NOT be null.

See [Appendix: Origination of cost data](#) section for examples of Provider, Publisher and Invoice Issuer values that can be used for various purchasing scenarios.

3.40.1. Column ID

ProviderName

3.40.2. Display Name

Provider

3.40.3. Description

The name of the entity that made the *resources* or *services* available for purchase.

3.40.4. Content Constraints

Constraint	Value
Column type	Dimension
Feature level	Mandatory
Allows nulls	False
Data type	String
Value format	<not specified>

3.40.5. Introduced (version)

0.5

3.41. Publisher

A Publisher is an entity that produces the [resources](#) or [services](#) that were purchased. It is commonly used for cost analysis and reporting scenarios.

The PublisherName column adheres to the following requirements:

- PublisherName MUST be present in a [FOCUS dataset](#).
- PublisherName MUST be of type String.
- PublisherName MUST conform to [StringHandling](#) requirements.
- PublisherName MUST NOT be null.

See [Appendix: Origination of cost data](#) section for examples of [Provider](#), Publisher and [Invoice Issuer](#) values that can be used for various purchasing scenarios.

3.41.1. Column ID

PublisherName

3.41.2. Display Name

Publisher

3.41.3. Description

The name of the entity that produced the *resources* or *services* that were purchased.

3.41.4. Content Constraints

Constraint	Value
Column type	Dimension
Feature level	Mandatory
Allows nulls	False
Data type	String

Constraint	Value
Value format	<not specified>

3.41.5. Introduced (version)

0.5

3.42. Region ID

A Region ID is a provider-assigned identifier for an isolated geographic area where a [resource](#) is provisioned or a [service](#) is provided. The region is commonly used for scenarios like analyzing cost and unit prices based on where *resources* are deployed.

The RegionId column adheres to the following requirements:

- RegionId MUST be present in a [FOCUS dataset](#) when the provider supports deploying resources or services within a region.
- RegionId MUST be of type String.
- RegionId MUST conform to [StringHandling](#) requirements.
- RegionId nullability is defined as follows:
 - RegionId MUST NOT be null when a *resource* or *service* is operated in or managed from a distinct region.
 - RegionId MAY be null when a *resource* or *service* is not operated in or managed from a distinct region.

3.42.1. Column ID

RegionId

3.42.2. Display Name

Region ID

3.42.3. Description

Provider-assigned identifier for an isolated geographic area where a *resource* is provisioned or a *service* is provided.

3.42.4. Content constraints

Constraint	Value
Column type	Dimension
Feature level	Conditional
Allows nulls	True
Data type	String
Value format	<not specified>

3.42.5. Introduced (version)

1.0

3.43. Region Name

Region Name is a provider-assigned display name for an isolated geographic area where a [resource](#) is provisioned or a [service](#) is provided. Region Name is commonly used for scenarios like analyzing cost and unit prices based on where *resources* are deployed.

The RegionName column adheres to the following requirements:

- RegionName MUST be present in a [FOCUS dataset](#) when the provider supports deploying resources or services within a region.
- RegionName MUST be of type String.
- RegionName MUST conform to [StringHandling](#) requirements.
- RegionName nullability is defined as follows:
 - RegionName MUST be null when [RegionId](#) is null.
 - RegionName MUST NOT be null when RegionId is not null.

3.43.1. Column ID

RegionName

3.43.2. Display Name

Region Name

3.43.3. Description

The name of an isolated geographic area where a *resource* is provisioned or a *service* is provided.

3.43.4. Content constraints

Constraint	Value
Column type	Dimension
Feature level	Conditional
Allows nulls	True
Data type	String
Value format	<not specified>

3.43.5. Introduced (version)

1.0

3.44. Resource ID

A Resource ID is an identifier assigned to a [resource](#) by the provider. The Resource ID is commonly used for cost reporting, analysis, and allocation scenarios.

The ResourceId column adheres to the following requirements:

- ResourceId MUST be present in a [FOCUS dataset](#) when the provider supports billing based on provisioned *resources*.
- ResourceId MUST be of type String.
- ResourceId MUST conform to [StringHandling](#) requirements.
- ResourceId nullability is defined as follows:
 - ResourceId MUST be null when a [charge](#) is not related to a *resource*.
 - ResourceId MUST NOT be null when a *charge* is related to a *resource*.
- When ResourceId is not null, ResourceId adheres to the following additional requirements:
 - ResourceId MUST be a unique identifier within the provider.
 - ResourceId SHOULD be a fully-qualified identifier.

3.44.1. Column ID

ResourceId

3.44.2. Display Name

Resource ID

3.44.3. Description

Identifier assigned to a *resource* by the provider.

3.44.4. Content Constraints

Constraint	Value
Column type	Dimension
Feature level	Conditional
Allows nulls	True
Data type	String
Value format	<not specified>

3.44.5. Introduced (version)

0.5

3.45. Resource Name

The Resource Name is a display name assigned to a [resource](#). It is commonly used for cost analysis, reporting, and allocation scenarios.

The ResourceName column adheres to the following requirements:

- ResourceName MUST be present in a [FOCUS dataset](#) when the provider supports billing based on provisioned resources.
- ResourceName MUST be of type String.
- ResourceName MUST conform to [StringHandling](#) requirements.
- ResourceName nullability is defined as follows:
 - ResourceName MUST be null when [ResourceId](#) is null or when the *resource* does not have an assigned display name.
 - ResourceName MUST NOT be null when ResourceId is not null and the *resource* has an assigned display name.
- ResourceName MUST NOT duplicate ResourceId when the *resource* is not provisioned interactively or only has a system-generated ResourceId.

3.45.1. Column ID

ResourceName

3.45.2. Display Name

Resource Name

3.45.3. Description

Display name assigned to a *resource*.

3.45.4. Content Constraints

Constraint	Value
Column type	Dimension
Feature level	Conditional
Allows nulls	True
Data type	String
Value format	<not specified>

3.45.5. Introduced (version)

0.5

3.46. Resource Type

Resource Type describes the kind of [resource](#) the [charge](#) applies to. A Resource Type is commonly used for scenarios like identifying cost changes in groups of similar *resources* and may include values like Virtual Machine, Data Warehouse, and Load Balancer.

The ResourceType column adheres to the following requirements:

- ResourceType MUST be present in a [FOCUS dataset](#) when the provider supports billing based on provisioned *resources* and supports assigning types to *resources*.
- ResourceType MUST be of type String.
- ResourceType MUST conform to [StringHandling](#) requirements.
- ResourceType nullability is defined as follows:
 - ResourceType MUST be null when [ResourceId](#) is null.
 - ResourceType MUST NOT be null when ResourceId is not null.

3.46.1. Column ID

ResourceType

3.46.2. Display Name

Resource Type

3.46.3. Description

The kind of *resource* the *charge* applies to.

3.46.4. Content Constraints

Constraint	Value
Column type	Dimension
Feature level	Conditional
Allows nulls	True
Data type	String
Value format	<not specified>

3.46.5. Introduced (version)

1.0-preview

3.47. Service Category

The Service Category is the highest-level classification of a [service](#) based on the core function of the *service*. Each *service* should have one and only one category that best aligns with its primary purpose. The Service Category is commonly used for scenarios like analyzing costs across providers and tracking the migration of workloads across fundamentally different architectures.

The ServiceCategory column adheres to the following requirements:

- ServiceCategory MUST be present in a [FOCUS dataset](#).
- ServiceCategory MUST be of type String.
- ServiceCategory MUST NOT be null.
- ServiceCategory MUST be one of the allowed values.

3.47.1. Column ID

ServiceCategory

3.47.2. Display Name

Service Category

3.47.3. Description

Highest-level classification of a *service* based on the core function of the *service*.

3.47.4. Content Constraints

Constraint	Value
Column type	Dimension
Feature level	Mandatory
Allows nulls	False
Data type	String
Value format	Allowed Values

Allowed values:

Service Category	Description
AI and Machine Learning	Artificial Intelligence and Machine Learning related technologies.
Analytics	Data processing, analytics, and visualization capabilities.
Business Applications	Business and productivity applications and services.
Compute	Virtual, containerized, serverless, or high-performance computing infrastructure and services.
Databases	Database platforms and services that allow for storage and querying of data.
Developer Tools	Software development and delivery tools and services.
Multicloud	Support for interworking of multiple cloud and/or on-premises environments.
Identity	Identity and access management services.
Integration	Services that allow applications to interact with one another.
Internet of Things	Development and management of IoT devices and networks.
Management and Governance	Management, logging, and observability of a customer's use of cloud.
Media	Media and entertainment streaming and processing services.
Migration	Moving applications and data to the cloud.
Mobile	Services enabling cloud applications to interact via mobile technologies.

Service Category	Description
Networking	Network connectivity and management.
Security	Security monitoring and compliance services.
Storage	Storage services for structured or unstructured data.
Web	Services enabling cloud applications to interact via the Internet.
Other	New or emerging services that do not align with an existing category.

3.47.5. Introduced (version)

0.5

3.48. Service Name

A [service](#) represents an offering that can be purchased from a provider (e.g., cloud virtual machine, SaaS database, professional services from a systems integrator). A *service* offering can include various types of usage or other [charges](#). For example, a cloud database *service* may include compute, storage, and networking *charges*.

The Service Name is a display name for the offering that was purchased. The Service Name is commonly used for scenarios like analyzing aggregate cost trends over time and filtering data to investigate anomalies.

The ServiceName column adheres to the following requirements:

- ServiceName MUST be present in a [FOCUS dataset](#).
- ServiceName MUST be of type String.
- ServiceName MUST conform to [StringHandling](#) requirements.
- ServiceName MUST NOT be null.
- The relationship between ServiceName and [ServiceCategory](#) is defined as follows:
 - ServiceName MUST have one and only one ServiceCategory that best aligns with its primary purpose, except when no suitable ServiceCategory is available.
 - ServiceName MUST be associated with the ServiceCategory "Other" when no suitable ServiceCategory is available.
- The relationship between ServiceName and [ServiceSubcategory](#) is defined as follows:
 - ServiceName SHOULD have one and only one ServiceSubcategory that best aligns with its primary purpose, except when no suitable ServiceSubcategory is available.
 - ServiceName SHOULD be associated with the ServiceSubcategory "Other" when no suitable ServiceSubcategory is available.

3.48.1. Column ID

ServiceName

3.48.2. Display Name

Service Name

3.48.3. Description

An offering that can be purchased from a provider (e.g., cloud virtual machine, SaaS database, professional *services* from a systems integrator).

3.48.4. Content Constraints

Constraint	Value
Column type	Dimension
Feature level	Mandatory
Allows nulls	False
Data type	String
Value format	<not specified>

3.48.5. Introduced (version)

0.5

3.49. Service Subcategory

The Service Subcategory is a secondary classification of the [Service Category](#) for a [service](#) based on its core function. The Service Subcategory (in conjunction with the Service Category) is commonly used for scenarios like analyzing spend and usage for specific workload types across providers and tracking the migration of workloads across fundamentally different architectures.

The ServiceSubcategory column adheres to the following requirements:

- ServiceSubcategory is RECOMMENDED to be present in a [FOCUS dataset](#).
- ServiceSubcategory MUST be of type String.
- ServiceSubcategory MUST NOT be null.
- ServiceSubcategory MUST be one of the allowed values.
- ServiceSubcategory MUST have one and only one parent ServiceCategory as specified in the allowed values below.

3.49.1. Column ID

ServiceSubcategory

3.49.2. Display Name

Service Subcategory

3.49.3. Description

Secondary classification of the Service Category for a *service* based on its core function.

3.49.4. Content Constraints

Constraint	Value
Column type	Dimension
Feature level	Recommended
Allows nulls	False
Data type	String
Value format	Allowed Values

Allowed values:

Service Category	Service Subcategory	Service Subcategory Description
AI and Machine Learning	AI Platforms	Unified solution that combines artificial intelligence and machine learning technologies.
AI and Machine Learning	Bots	Automated performance of tasks such as customer service, data collection, and content moderation.
AI and Machine Learning	Generative AI	Creation of content like text, images, and music by learning patterns from existing data.
AI and Machine Learning	Machine Learning	Creation, training, and deployment of statistical algorithms that learn from and perform tasks based on data.
AI and Machine Learning	Natural Language Processing	Generation of human language, handling tasks like translation, sentiment analysis, and text summarization.
AI and Machine Learning	Other (AI and Machine Learning)	AI and Machine Learning services that do not fall into one of the defined subcategories.
Analytics	Analytics Platforms	Unified solution that combines technologies across the entire analytics lifecycle.
Analytics	Business Intelligence	Semantic models, dashboards, reports, and data visualizations to track performance and identify trends.
Analytics	Data Processing	Integration and transformation tasks to prepare data for analysis.
Analytics	Search	Discovery of information by indexing and retrieving data from various sources.
Analytics	Streaming Analytics	Real-time data stream processes to detect patterns, trends, and anomalies as they occur.
Analytics	Other (Analytics)	Analytics services that do not fall into one of the defined subcategories.
Business Applications	Productivity and Collaboration	Tools that facilitate individuals managing tasks and working together.
Business Applications	Other (Business Applications)	Business Applications services that do not fall into one of the defined subcategories.
Compute	Containers	Management and orchestration of containerized compute platforms.
Compute	End User Computing	Virtualized desktop infrastructure and device / endpoint management.
Compute	Quantum Compute	Resources and simulators that leverage the principles of quantum mechanics.

Service Category	Service Subcategory	Service Subcategory Description
Compute	Serverless Compute	Enablement of compute capabilities without provisioning or managing servers.
Compute	Virtual Machines	Computing environments ranging from hosts with abstracted operating systems to bare-metal servers.
Compute	Other (Compute)	Compute services that do not fall into one of the defined subcategories.
Databases	Caching	Low-latency and high-throughput access to frequently accessed data.
Databases	Data Warehouses	Big data storage and querying capabilities.
Databases	Ledger Databases	Immutable and transparent databases to record tamper-proof and cryptographically secure transactions.
Databases	NoSQL Databases	Unstructured or semi-structured data storage and querying capabilities.
Databases	Relational Databases	Structured data storage and querying capabilities.
Databases	Time Series Databases	Time-stamped data storage and querying capabilities.
Databases	Other (Databases)	Database services that do not fall into one of the defined subcategories.
Developer Tools	Developer Platforms	Unified solution that combines technologies across multiple areas of the software development lifecycle.
Developer Tools	Continuous Integration and Deployment	CI/CD tools and services that support building and deploying code for software and systems.
Developer Tools	Development Environments	Tools and services that support authoring code for software and systems.
Developer Tools	Source Code Management	Tools and services that support version control of code for software and systems.
Developer Tools	Quality Assurance	Tools and services that support testing code for software and systems.
Developer Tools	Other (Developer Tools)	Developer Tools services that do not fall into one of the defined subcategories.
Identity	Identity and Access Management	Technologies that ensure users have appropriate access to resources.
Identity	Other (Identity)	Identity services that do not fall into one of the defined subcategories.
Integration	API Management	Creation, publishing, and management of application programming interfaces.
Integration	Messaging	Asynchronous communication between distributed applications.
Integration	Workflow Orchestration	Design, execution, and management of business processes and workflows.
Integration	Other (Integration)	Integration services that do not fall into one of the defined subcategories.
Internet of Things	IoT Analytics	Examination of data collected from IoT devices.
Internet of Things	IoT Platforms	Unified solution that combines IoT data collection, processing, visualization, and device management.

Service Category	Service Subcategory	Service Subcategory Description
Internet of Things	Other (Internet of Things)	Internet of Things (IoT) services that do not fall into one of the defined subcategories.
Management and Governance	Architecture	Planning, design, and construction of software systems.
Management and Governance	Compliance	Adherence to regulatory standards and industry best practices.
Management and Governance	Cost Management	Monitoring and controlling expenses of systems and services.
Management and Governance	Data Governance	Management of the availability, usability, integrity, and security of data.
Management and Governance	Disaster Recovery	Plans and procedures that ensure systems and services can recover from disruptions.
Management and Governance	Endpoint Management	Tools that configure and secure access to devices.
Management and Governance	Observability	Monitoring, logging, and tracing of data to track the performance and health of systems.
Management and Governance	Support	Assistance and expertise supplied by providers.
Management and Governance	Other (Management and Governance)	Management and governance services that do not fall into one of the defined subcategories.
Media	Content Creation	Production of media content.
Media	Gaming	Development and delivery of gaming services.
Media	Media Streaming	Multimedia delivered and rendered in real-time on devices.
Media	Mixed Reality	Technologies that blend real-world and computer-generated environments.
Media	Other (Media)	Media services that do not fall into one of the defined subcategories.
Migration	Data Migration	Movement of stored data from one location to another.
Migration	Resource Migration	Movement of resources from one location to another.
Migration	Other (Migration)	Migration services that do not fall into one of the defined subcategories.
Mobile	Other (Mobile)	All Mobile services.
Multicloud	Multicloud Integration	Environments that facilitate consumption of services from multiple cloud providers.
Multicloud	Other (Multicloud)	Multicloud services that do not fall into one of the defined subcategories.
Networking	Application Networking	Distribution of incoming network traffic across application-based workloads.
Networking	Content Delivery	Distribution of digital content using a network of servers (CDNs).

Service Category	Service Subcategory	Service Subcategory Description
Networking	Network Connectivity	Facilitates communication between networks or network segments.
Networking	Network Infrastructure	Configuration, monitoring, and troubleshooting of network devices.
Networking	Network Routing	Services that select paths for traffic within or across networks.
Networking	Network Security	Protection from unauthorized network access and cyber threats using firewalls and anti-malware tools.
Networking	Other (Networking)	Networking services that do not fall into one of the defined subcategories.
Security	Secret Management	Information used to authenticate users and systems, including secrets, certificates, tokens, and other keys.
Security	Security Posture Management	Tools that help organizations configure, monitor, and improve system security.
Security	Threat Detection and Response	Collect and analyze security data to identify and respond to potential security threats and vulnerabilities.
Security	Other (Security)	Security services that do not fall into one of the defined subcategories.
Storage	Backup Storage	Secondary storage to protect against data loss.
Storage	Block Storage	High performance, low latency storage that provides random access.
Storage	File Storage	Scalable, sharable storage for file-based data.
Storage	Object Storage	Highly available, durable storage for unstructured data.
Storage	Storage Platforms	Unified solution that supports multiple storage types.
Storage	Other (Storage)	Storage services that do not fall into one of the defined subcategories.
Web	Application Platforms	Integrated environments that run web applications.
Web	Other (Web)	Web services that do not fall into one of the defined subcategories.
Other	Other (Other)	Services that do not fall into one of the defined categories.

3.49.5. Introduced (version)

1.1

3.50. SKU ID

A SKU ID is a provider-specified unique identifier that represents a specific [SKU](#). *SKUs* are quantifiable goods or service offerings in a [FOCUS dataset](#) that represent specific functionality and technical specifications. Examples of *SKUs* include but are not limited to:

- A product license that is purchased or subscribed to.
- Usage of a deployed resource from direct user interaction (e.g., request count).
- Usage by a deployed resource based on the resource's configuration (e.g., running hours, storage space).

Each SKU ID represents a unique set of features that can be sold at different price points or [SKU](#)

[Prices](#). SKU ID is consistent across all pricing variations, which may differ based on multiple factors beyond the common functionality and technical specifications. Examples include but are not limited to:

- Date the [charge](#) was incurred.
- Pricing tiers (e.g., free tier or volume-based tiers).
- Commitment discount pricing terms (e.g., 1 year, 3 years).
- Negotiated discounts or other contractual terms or conditions.

SKU ID should be consistent across pricing variations of a good or service to facilitate price comparisons for the same functionality, like where the functionality is provided or how it's paid for. SKU ID can be referenced on a catalog or [price list](#) published by a provider to look up detailed information about the *SKU*. The composition of the properties associated with the SKU ID may differ across providers. SKU ID is commonly used for analyzing and comparing costs for the same SKU across different price details (e.g., term, tier, location).

The Skuld column adheres to the following requirements:

- Skuld MUST be present in a [FOCUS dataset](#) when the provider supports unit pricing concepts and publishes price lists, publicly or as part of contracting.
- Skuld MUST be of type String.
- Skuld MUST conform to [StringHandling](#) requirements.
- Skuld nullability is defined as follows:
 - Skuld MUST be null when [ChargeCategory](#) is "Tax".
 - Skuld MUST NOT be null when ChargeCategory is "Usage" or "Purchase" and [ChargeClass](#) is not "Correction".
 - Skuld MAY be null in all other cases.
- Skuld for a given *SKU* adheres to the following additional requirements:
 - Skuld MUST remain consistent across [billing accounts](#) or contracts.
 - Skuld MUST remain consistent across [PricingCategory](#) values.
 - Skuld MUST remain consistent regardless of any other factors that might impact the price but do not affect the functionality of the *SKU*.
- Skuld MUST be associated with a given [resource](#) or [service](#) when ChargeCategory is "Usage" or "Purchase".
- Skuld MAY equal [SkuPriceld](#).

3.50.1. Column ID

Skuld

3.50.2. Display Name

SKU ID

3.50.3. Description

Provider-specified unique identifier that represents a specific *SKU* (e.g., a quantifiable good or service offering).

3.50.4. Content constraints

Constraint	Value
------------	-------

Constraint	Value
Column type	Dimension
Feature level	Conditional
Allows nulls	True
Data type	String
Value format	<not specified>

3.50.5. Introduced (version)

1.0-preview

3.51. SKU Meter

SKU Meter describes the functionality being metered or measured by a particular SKU in a [charge](#).

Providers often have billing models in which multiple SKUs exist for a given service to describe and bill for different functionalities for that service. For example, an object storage service may have separate SKUs for functionalities such as object storage, API requests, data transfer, encryption, and object management. This field helps practitioners understand which functionalities are being metered by the different SKUs that appear in a [FOCUS dataset](#).

The SkuMeter column adheres to the following requirements:

- SkuMeter MUST be present in a *FOCUS dataset* when the provider supports unit pricing concepts and publishes [price lists](#), publicly or as part of contracting.
- SkuMeter MUST be of type String.
- SkuMeter MUST conform to [StringHandling](#) requirements.
- SkuMeter nullability is defined as follows:
 - SkuMeter MUST be null when [Skuld](#) is null.
 - SkuMeter SHOULD NOT be null when Skuld is not null.
- SkuMeter SHOULD remain consistent over time for a given Skuld.

3.51.1. Examples

Compute Usage, Block Volume Usage, Data Transfer, API Requests

3.51.2. Column ID

SkuMeter

3.51.3. Display Name

SKU Meter

3.51.4. Description

Describes the functionality being metered or measured by a particular SKU in a *charge*.

3.51.5. Content Constraints

Constraint	Value
Column type	Dimension
Feature level	Conditional
Allows nulls	True
Data type	String
Value format	<not specified>

3.51.6. Introduced (version)

1.1

3.52. SKU Price Details

SKU Price Details represent a list of [SKU Price](#) properties (key-value pairs) associated with a specific [SKU Price ID](#). These properties include qualitative and quantitative properties of a [SKU](#) (e.g., functionality and technical specifications), along with core stable pricing properties (e.g., pricing terms, tiers, etc.), excluding dynamic or negotiable pricing elements such as unit price amounts, currency (and related exchange rates), temporal validity (e.g., effective dates), and contract- or negotiation-specific factors (e.g., contract or account identifiers, and negotiable discounts).

The composition of properties associated with a specific *SKU Price* may differ across providers and across *SKUs* within the same provider. However, the exclusion of dynamic or negotiable pricing properties should ensure that all [charges](#) with the same SKU Price ID share the same SKU Price Details, i.e., that SKU Price Details remains consistent across different [billing periods](#) and [billing accounts](#) within a provider.

SKU Price Details helps practitioners understand and distinguish *SKU Prices*, each identified by a SKU Price ID and associated with a used or purchased [resource](#) or [service](#). It can also help determine the quantity of units for a property when it holds a numeric value (e.g., CoreCount), even when its unit differs from the one in which the *SKU* is priced and charged, thus supporting FinOps capabilities like unit economics. Additionally, the SKU Price Details may be used to analyze costs based on pricing properties such as terms and tiers.

The SkuPriceDetails column adheres to the following requirements:

- SkuPriceDetails MUST be present in a [FOCUS dataset](#) when the provider supports unit pricing concepts and publishes [price lists](#), publicly or as part of contracting.
- SkuPriceDetails MUST conform to [KeyValueFormat](#) requirements.
- SkuPriceDetails property keys SHOULD conform to [PascalCase](#) format.
- SkuPriceDetails nullability is defined as follows:
 - SkuPriceDetails MUST be null when SkuPricId is null.
 - SkuPriceDetails MAY be null when SkuPricId is not null.
- When SkuPriceDetails is not null, SkuPriceDetails adheres to the following additional requirements:
 - SkuPriceDetails MUST be associated with a given SkuPricId.
 - SkuPriceDetails MUST NOT include properties that are not applicable to the corresponding SkuPricId.
 - SkuPriceDetails SHOULD include all FOCUS-defined SKU Price properties listed below that are applicable to the corresponding SkuPricId.
 - SkuPriceDetails MUST include the FOCUS-defined SKU Price property when an equivalent property is included as a Provider-defined property.

- SkuPriceDetails MAY include properties that are already captured in other dedicated columns.
- SkuPriceDetails properties for a given SkuPriceId adhere to the following additional requirements:
 - Existing SkuPriceDetails properties SHOULD remain consistent over time.
 - Existing SkuPriceDetails properties SHOULD NOT be removed.
 - Additional SkuPriceDetails properties MAY be added over time.
- Property key SHOULD remain consistent across comparable *SKUs* having that property, and the values for this key SHOULD remain in a consistent format.
- Property key MUST begin with the string "x_" unless it is a FOCUS-defined property.
- Property value MUST represent the value for a single [PricingUnit](#) when the property holds a numeric value.
- FOCUS-defined SKU Price properties adhere to the following additional requirements:
 - Property key MUST match the spelling and casing specified for the FOCUS-defined property.
 - Property value MUST be of the type specified for that property.
 - Property value MUST represent the value for a single PricingUnit, denominated in the unit of measure specified for that property when the property holds a numeric value.

3.52.1. Examples

```
{
  "StorageClass": "Archive",
  "CoreCount": 4,
  "x_PremiumProcessing": true,
}
```

3.52.2. Column ID

SkuPriceDetails

3.52.3. Display Name

SKU Price Details

3.52.4. Description

A set of properties of a SKU Price ID which are meaningful and common to all instances of that SKU Price ID.

3.52.5. Content Constraints

Constraint	Value
Column type	Dimension
Feature level	Conditional
Allows nulls	True
Data type	JSON

Constraint	Value
Value format	KeyValueFormat

3.52.5.1. FOCUS-Defined Properties

The following keys should be used when applicable to facilitate cross-SKU and cross-provider queries for the same conceptual property. Focus-defined keys will appear in the list below and Provider-defined keys will be prefixed with "x_" to make them easy to identify as well as prevent collisions.

Key	Description	Data Type	Unit of Measure (numeric) or example values (string)
CoreCount	Number of physical or virtual CPUs available ¹	Numeric	Measure: Quantity of Cores
DiskMaxIops	Storage maximum sustained input/output operations per second ¹	Numeric	Measure: Input/Output Operations per Second (IOPS)
DiskSpace	Storage capacity available	Numeric	Measure: Gibibytes (GiB)
DiskType	Kind of disk used	String	Examples: "SSD", "HDD", "NVMe"
GpuCount	Number of GPUs available	Numeric	Measure: Quantity of GPUs
InstanceType	Common name of the instance including size, shape, series, etc.	String	Examples: "m5d.2xlarge", "NC24rs_v3", "P50"
InstanceSeries	Common name for the series and/or generation of the instance	String	Examples: "M5", "DadV5", "N2D"
MemorySize	RAM allocated for processing	Numeric	Measure: Gibibytes (GiB ²)
NetworkMaxIops	Network maximum sustained input/output operations per second ¹	Numeric	Measure: Input/Output Operations per Second (IOPS)
NetworkMaxThroughput	Network maximum sustained throughput for data transfer ¹	Numeric	Measure: Megabits per second (Mbps)
OperatingSystem	Operating system family ³	String	Examples: "Linux", "MacOS", "Windows"
Redundancy	Level of redundancy offered by the SKU	String	Examples: "Local", "Zonal", "Global"
StorageClass	Class or tier of storage provided	String	Examples: "Hot", "Archive", "Nearline"

Notes

¹ In the case of "burstable" SKUs offering variable levels of performance, the baseline or guaranteed value should be used.

² Memory manufacturers still commonly uses "GB" to refer to 2³⁰ bytes, which is known as GiB in other contexts.

³ This is the operating system family of the SKU, if it's included with the SKU or the SKU only supports one type of operating system.

3.52.6. Introduced (version)

1.1

3.53. SKU Price ID

SKU Price ID is a provider-specified unique identifier that represents a specific [SKU Price](#) associated with a [resource](#) or [service](#) used or purchased. It serves as a key reference for a *SKU Price* in a [price list](#) published by a provider, allowing practitioners to look up detailed information about the *SKU Price*.

The composition of properties associated with the SKU Price ID may differ across providers and across *SKUs* within the same provider. However, the exclusion of dynamic or negotiable pricing properties, such as unit price amount, currency (and related exchange rates), temporal validity (e.g., effective dates), and contract- or negotiation-specific elements (e.g., contract or account identifiers, and negotiable discounts), ensures that the SKU Price ID remains consistent across different billing periods and billing accounts within a provider. This consistency enables efficient filtering of [charges](#) to track price fluctuations (e.g., changes in unit price amounts) over time and across billing accounts, for both list and contracted unit prices. Additionally, the SKU Price ID is commonly used to analyze costs based on pricing properties such as terms and tiers.

The `SkuPriceId` column adheres to the following requirements:

- `SkuPriceId` MUST be present in a [FOCUS dataset](#) when the provider supports unit pricing concepts and publishes *price lists*, publicly or as part of contracting.
- `SkuPriceId` MUST be of type String.
- `SkuPriceId` MUST conform to [String Handling](#) requirements.
- `SkuPriceId` nullability is defined as follows:
 - `SkuPriceId` MUST be null when [ChargeCategory](#) is "Tax".
 - `SkuPriceId` MUST NOT be null when [ChargeCategory](#) is "Usage" or "Purchase" and [ChargeClass](#) is not "Correction".
 - `SkuPriceId` MAY be null in all other cases.
- When `SkuPriceId` is not null, `SkuPriceId` adheres to the following additional requirements:
 - `SkuPriceId` MUST have one and only one parent [Skuld](#).
 - `SkuPriceId` MUST remain consistent over time.
 - `SkuPriceId` MUST remain consistent across [billing accounts](#) or contracts.
 - `SkuPriceId` MAY equal `Skuld`.
 - `SkuPriceId` MUST be associated with a given [resource](#) or [service](#) when [ChargeCategory](#) is "Usage" or "Purchase".
 - `SkuPriceId` MUST reference a *SKU Price* in a provider-supplied *price list*, enabling the lookup of detailed information about the *SKU Price*.
 - `SkuPriceId` MUST support the lookup of the [ListUnitPrice](#) when the provider publishes unit prices exclusive of discounts.
 - `SkuPriceId` MUST support the verification of the given [ContractedUnitPrice](#) when the provider supports negotiated pricing concepts.

See [Examples: Commitment Discount Flexibility](#) for more details around *commitment discount flexibility*.

3.53.1. Column ID

`SkuPriceId`

3.53.2. Display Name

SKU Price ID

3.53.3. Description

A provider-specified unique identifier that represents a specific *SKU Price* associated with a *resource* or *service* used or purchased.

3.53.4. Content constraints

Constraint	Value
Column type	Dimension
Feature level	Conditional
Allows nulls	True
Data type	String
Value format	<not specified>

3.53.5. Introduced (version)

1.0-preview

3.54. Sub Account ID

A Sub Account ID is a provider-assigned identifier assigned to a [sub account](#). Sub Account ID is commonly used for scenarios like grouping based on organizational constructs, access management needs, and cost allocation strategies.

The SubAccountId column adheres to the following requirements:

- SubAccountId MUST be present in a [FOCUS dataset](#) when the provider supports a *sub account* construct.
- SubAccountId MUST be of type String.
- SubAccountId MUST conform to [StringHandling](#) requirements.
- SubAccountId nullability is defined as follows:
 - SubAccountId MUST be null when a [charge](#) is not related to a *sub account*.
 - SubAccountId MUST NOT be null when a *charge* is related to a *sub account*.

See [Appendix: Grouping constructs for resources or services](#) for details and examples of the different grouping constructs supported by FOCUS.

3.54.1. Column ID

SubAccountId

3.54.2. Display Name

3.54.3. Description

An ID assigned to a grouping of [resources](#) or [services](#), often used to manage access and/or cost.

3.54.4. Content constraints

Constraint	Value
Column type	Dimension
Feature level	Conditional
Allows nulls	True
Data type	String
Value format	<not specified>

3.54.5. Introduced (version)

0.5

3.55. Sub Account Name

A Sub Account Name is a display name assigned to a [sub account](#). Sub account Name is commonly used for scenarios like grouping based on organizational constructs, access management needs, and cost allocation strategies.

The SubAccountName column adheres to the following requirements:

- SubAccountName MUST be present in a [FOCUS dataset](#) when the provider supports a *sub account* construct.
- SubAccountName MUST be of type String.
- SubAccountName MUST conform to [StringHandling](#) requirements.
- SubAccountName nullability is defined as follows:
 - SubAccountName MUST be null when [SubAccountId](#) is null.
 - SubAccountName MUST NOT be null when SubAccountId is not null.

See [Appendix: Grouping constructs for resources or services](#) for details and examples of the different grouping constructs supported by FOCUS.

3.55.1. Column ID

SubAccountName

3.55.2. Display Name

Sub Account Name

3.55.3. Description

A name assigned to a grouping of [resources](#) or [services](#), often used to manage access and/or cost.

3.55.4. Content constraints

Constraint	Value
Column type	Dimension
Feature level	Conditional
Allows nulls	True
Data type	String
Value format	<not specified>

3.55.5. Introduced (version)

0.5

3.56. Sub Account Type

Sub Account Type is a provider-assigned name to identify the type of [sub account](#). Sub Account Type is a readable display name and not a code. Sub Account Type is commonly used for scenarios like mapping FOCUS and provider constructs, summarizing costs across providers, or invoicing and chargeback.

The SubAccountType column adheres to the following requirements:

- SubAccountType MUST be present in a [FOCUS dataset](#) when the provider supports more than one possible SubAccountType value.
- SubAccountType MUST be of type String.
- SubAccountType MUST conform to [StringHandling](#) requirements.
- SubAccountType nullability is defined as follows:
 - SubAccountType MUST be null when [SubAccountId](#) is null.
 - SubAccountType MUST NOT be null when SubAccountId is not null.
- SubAccountType MUST be a consistent, readable display value.

3.56.1. Column ID

SubAccountType

3.56.2. Display Name

Sub Account Type

3.56.3. Description

A provider-assigned name to identify the type of *sub account*.

3.56.4. Content Constraints

Constraint	Value
Column type	Dimension
Feature level	Conditional
Allows nulls	True
Data type	String
Value format	<not specified>

3.56.5. Introduced (version)

1.2

3.57. Tags

The Tags column represents the set of [tags](#) assigned to [tag sources](#) that also account for potential provider-defined or user-defined tag evaluations. Tags are commonly used for scenarios like adding business context to cost and usage data to identify and accurately allocate [charges](#). Tags may also be referred to by providers using other terms such as labels.

A tag becomes [finalized](#) when a single value is selected from a set of possible tag values assigned to the tag key. When supported by a provider, this can occur when a tag value is set by provider-defined or user-defined rules.

The Tags column adheres to the following requirements:

- Tags MUST be present in a [FOCUS dataset](#) when the provider supports setting user or provider-defined tags.
- Tags MUST conform to [KeyValueFormat](#) requirements.
- Tags MAY be null.
- When Tags is not null, Tags adheres to the following additional requirements:
 - Tags MUST include all user-defined and provider-defined tags.
 - Tags MUST only include finalized tags.
 - Tags SHOULD include tag keys with corresponding non-null values for a given [resource](#).
 - Tags MAY include tag keys with a null value for a given [resource](#) depending on the provider's tag finalization process.
 - Tag keys that do not support corresponding values, MUST have a corresponding true (boolean) value set.
 - Provider SHOULD publish tag finalization methods and semantics within their respective documentation.
 - Provider MUST NOT alter tag values unless applying true (boolean) to valueless tags.
- Provider-defined tags adhere to the following additional requirements:
 - Provider-defined tag keys MUST be prefixed with a predetermined, provider-specified tag key prefix that is unique to each corresponding provider-specified tag scheme.
 - Provider SHOULD publish all provider-specified tag key prefixes within their respective documentation.
- User-defined tags adhere to the following additional requirements:
 - Provider MUST prefix all but one user-defined tag scheme with a predetermined, provider-specified tag key prefix that is unique to each corresponding user-defined tag scheme when the provider has more than one user-defined tag scheme.
 - Provider MUST NOT prefix tag keys when the provider has only one user-defined tag scheme.
 - Provider MUST NOT allow reserved tag key prefixes to be used as prefixes for any user-

defined tag keys within a prefixless user-defined tag scheme.

3.57.1. Provider-Defined vs. User-Defined Tags

This example illustrates various tags produced from multiple user-defined and provider-defined tag schemes. The first three tags illustrate examples from three different, user-defined tag schemes. The provider predetermined that 1 user-defined tag scheme (i.e., "foo": "bar") does not have a prepended prefix, but the remaining two user-defined tag schemes (i.e., "userDefinedTagScheme2/foo": "bar", "userDefinedTagScheme3/foo": true) do have provider-defined and reserved prefixes. Additionally, the third tag is produced from a valueless, user-defined tag scheme, so the provider also applies true as its default value.

The last two tags illustrate examples from two different, provider-defined tag schemes. Since all provider-defined tag schemes require a prefix, the provider has prepended predefined and reserved prefixes (providerDefinedTagScheme1/, providerDefinedTagScheme2/) to each tag.

```
{
  "foo": "bar",
  "userDefinedTagScheme2/foo": "bar",
  "userDefinedTagScheme3/foo": true,
  "providerDefinedTagScheme1/foo": "bar",
  "providerDefinedTagScheme2/foo": "bar"
}
```

3.57.2. Finalized Tags

Within a provider, tag keys may be associated with multiple values, and potentially defined at different levels within the provider, such as accounts, folders, [resource](#) and other *resource* grouping constructs. When finalizing, *providers* must reduce these multiple levels of definition to a single value where each key is associated with exactly one value. The method by which this is done and the semantics are up to each provider but must be documented within their respective documentation.

As an example, let's assume 1 [sub account](#) exists with 1 virtual machine with the following details, and tag inheritance favors Resources over *Sub Accounts*.

- Sub Account
 - id: *my-sub-account*
 - user-defined tags: *team:ops*, *env:prod*
- Virtual Machine
 - id: *my-vm*
 - user-defined tags: *team:web*

The table below represents a finalized dataset with these *resources*. It also shows the finalized state after all resource-oriented, tag inheritance rules are processed.

ResourceType	ResourceId	Tags
Sub Account	my-sub-account	{ "team": "ops", "env": "prod" }
Virtual Machine	my-vm	{ "team": "web", "env": "prod" }

Because the Virtual Machine Resource did not have an env tag, it inherited tag, env:prod (italicized), from its parent *sub account*. Conversely, because the Virtual Machine Resource already has a team tag (team:web), it did not inherit team:ops from its parent *sub account*.

3.57.3. Column ID

Tags

3.57.4. Display Name

Tags

3.57.5. Description

The set of tags assigned to *tag sources* that account for potential provider-defined or user-defined tag evaluations.

3.57.6. Content Constraints

Constraint	Value
Column type	Dimension
Feature level	Conditional
Allows nulls	True
Data type	JSON
Value format	Key-Value Format

3.57.7. Introduced (version)

1.0-preview

4. Attributes

Attributes are requirements that apply across a [FOCUS dataset](#) instead of an individual column level. Requirements on data content can include naming conventions, data types, formatting standardizations, etc. Attributes may introduce high-level requirements for data granularity, recency, frequency, etc. Requirements defined in attributes are necessary for servicing [FinOps capabilities](#) accurately using a standard set of instructions regardless of the origin of the data.

4.1. Column Handling

A [FOCUS dataset](#) consists of a set of columns that convey information about the charges incurred with a provider. Each column describes an aspect of the charge, including but not limited to:

- Who is responsible for incurring or delivering the service.
- What the charge is for.
- When the charge was incurred.
- Where the service was delivered.
- Why the charge was incurred for a specific price.
- How much the charge is and how that cost is calculated.

While FOCUS establishes the core structure and standardizes columns for consistent reporting of cost and usage data, the diverse and evolving landscape of providers and service offerings may require providers and data generators to include supplemental columns in the FOCUS dataset. These additional columns may enable deeper analysis and provide more detailed descriptions of usage that may not be fully captured by standard FOCUS dataset columns.

In such cases, providers and data generators are responsible for ensuring that their usage and cost data is accurately and comprehensively represented by including necessary supplemental columns without duplicating data in FOCUS columns. Rows in a FOCUS dataset may be aggregated or split differently than non-FOCUS datasets to align with FOCUS requirements (e.g., Discount Handling), while enriching the dataset, providers and data generators must maintain the integrity of FOCUS-defined dimensions and metrics. When performing these transformations, providers and data generators must ensure the accuracy of all dimensions and metrics, particularly summable values such as costs and quantities.

Columns within FOCUS include an ID and a display name. Column IDs are used in files and database tables and display names can be used in report output and other descriptive content, like documentation. Column IDs provided in a *FOCUS dataset* follow consistent naming and ordering conventions for FinOps practitioners who consume the data for analysis, reporting, and other use cases.

All columns defined in the FOCUS specification MUST follow the naming and ordering requirements listed below.

4.1.1. Attribute ID

ColumnHandling

4.1.2. Attribute Name

Column Handling

4.1.3. Description

Naming and ordering convention for columns appearing in a *FOCUS dataset*.

4.1.4. Requirements

4.1.4.1. Column Names

- All columns defined by FOCUS MUST follow the following rules:
 - Column IDs MUST use [Pascal case](#).
 - Column IDs MUST NOT use abbreviations.
 - Column IDs MUST be alphanumeric with no special characters.
 - Column IDs SHOULD NOT use acronyms.
 - Column IDs SHOULD NOT exceed 50 characters to accommodate column length restrictions of various data repositories.
 - Columns that have an ID and a Name MUST have the Id or Name suffix in the Column ID.
 - Column display names MAY avoid the Name suffix if there are no other columns with the same name prefix.

- Columns with the Category suffix MUST be normalized.
- Custom (e.g., provider-defined) columns that are not defined by FOCUS but included in a *FOCUS dataset* MUST follow the following rules:
 - Custom columns MUST be prefixed with a consistent x_ prefix to identify them as external, custom columns and distinguish them from FOCUS columns to avoid conflicts in future releases.
 - Custom columns SHOULD follow the same rules listed above for FOCUS columns.

4.1.4.2. Column Order

- All FOCUS columns SHOULD be first in the provided dataset.
- Custom columns SHOULD be listed after all FOCUS columns and SHOULD NOT be intermixed.
- Columns MAY be sorted alphabetically, but custom columns SHOULD be after all FOCUS columns.

4.1.5. Exceptions

- Identifiers will use the "Id" abbreviation since this is a standard pattern across the industry.
- Product offerings that incur charges will use the "Sku" abbreviation because it is a well-understood term both within and outside the industry.

4.1.6. Introduced (version)

0.5

4.2. Currency Format

Columns that contain currency information in cost data following a consistent format reduce friction for FinOps practitioners who consume the data for analysis, reporting, and other use cases.

A currency may be one of the following currency types:

- National currency (e.g., USD, EUR).
- Virtual currency (e.g., tokens, credits).

All columns capturing a currency value, defined in the FOCUS specification, MUST follow the requirements listed below. Custom currency-related columns SHOULD also follow the same formatting requirements.

4.2.1. Attribute ID

CurrencyFormat

4.2.2. Attribute Name

Currency Format

4.2.3. Description

Formatting for currency columns appearing in a [FOCUS dataset](#).

4.2.4. Requirements

- Currency-related columns MUST be represented as a three-letter alphabetic code as dictated in the governing document [ISO 4217:2015](#) when the value is presented in national currency (e.g., USD, EUR).
- Currency-related columns MUST conform to [StringHandling](#) requirements when the value is presented in virtual currency (e.g., credits, tokens).

4.2.5. Exceptions

None

4.2.6. Introduced (version)

0.5

4.3. Date/Time Format

Columns that provide date and time information conforming to specified rules and formatting requirements ensure clarity, accuracy, and ease of interpretation for both humans and systems.

All columns capturing a date/time value, defined in the FOCUS specification, MUST follow the formatting requirements listed below. Custom date/time-related columns SHOULD also follow the same formatting requirements.

4.3.1. Attribute ID

DateTimeFormat

4.3.2. Attribute Name

Date/Time Format

4.3.3. Description

Rules and formatting requirements for date/time-related columns appearing in a [FOCUS dataset](#).

4.3.4. Requirements

- Date/time values MUST be in UTC (Coordinated Universal Time) to avoid ambiguity and ensure

consistency across different time zones.

- Date/time values format MUST be aligned with ISO 8601 standard, which provides a globally recognized format for representing dates and times (see [ISO 8601-1:2019](#) governing document for details).
- Values providing information about a specific moment in time MUST be represented in the extended ISO 8601 format with UTC offset ('YYYY-MM-DDTHH:mm:ssZ') and conform to the following guidelines:
 - Include the date and time components, separated with the letter 'T'
 - Use two-digit hours (HH), minutes (mm), and seconds (ss).
 - End with the 'Z' indicator to denote UTC (Coordinated Universal Time)

4.3.5. Exceptions

None

4.3.6. Introduced (version)

0.5

4.4. Discount Handling

A discount is a pricing construct where providers offer a reduced price for [services](#). Providers may have many types of discounts, including but not limited to commercially [negotiated discounts](#), [commitment discounts](#) when you agree to a certain amount of usage or spend, and bundled discounts where you receive free or discounted usage of one product or *service* based on the usage of another. Discount Handling is commonly used in scenarios like verifying discounts were applied and calculating cost savings.

Some discount offers can be purchased from a provider to get reduced prices. The most common example is a *commitment discount*, where you "purchase" a commitment to use or spend a specific amount within a period. When a commitment isn't fully utilized, the unused amount reduces the potential savings from the discount and can even result in paying higher costs than without the discount. Due to this risk, unused commitment amounts need to be clearly identifiable at a granular level. To facilitate this, unused commitments are recorded with a separate row for each charge period where the commitment was not fully utilized. To show the impact of purchased discounts on each discounted row, discount purchases need the purchase amount to be amortized over the [term](#) the discount is applied to (e.g., 1 year) with each [charge period](#) split and applied to each row that received the discount.

Amortization is a process used to break down and spread purchase costs over a period of time or *term* of use. When a purchase is applicable to resources, like *commitment discounts*, the amortized cost of a resource takes the initial payment and *term* into account and distributes it out based on the resource's usage, attributing the prorated cost for each unit of billing. Amortization enables users of billing data to distribute purchase charges to the appropriate audience in support of cost allocation efforts. Discount Handling for purchased commitments is commonly used for scenarios like calculating utilization and implementing chargeback for the purchase amount.

While providers may use different terms to describe discounts, FOCUS identifies a discount as being a reduced price applied directly to a row. Any price or cost reductions that are awarded after the fact are identified as a "Credit" Charge Category. One example might be when a provider offers a reduced rate after passing a certain threshold of usage or spend.

All rows defined in FOCUS MUST follow the discount handling requirements listed below.

4.4.1. Attribute ID

DiscountHandling

4.4.2. Attribute Name

Discount Handling

4.4.3. Description

Indicates how to include and apply discounts to usage charges or rows in a FOCUS dataset.

4.4.4. Requirements

- All applicable discounts SHOULD be applied to each row they pertain to and SHOULD NOT be negated in a separate row.
- All discounts applied to a row MUST apply to the entire charge.
 - Multiple discounts MAY apply to a row, but they MUST apply to the entire charge covered by that row.
 - If a discount only applies to a portion of a charge, then the discounted portion of the charge MUST be split into a separate row.
 - Each discount MUST be identifiable using existing FOCUS columns.
 - Rows with a *commitment discount* applied to them MUST include a CommitmentDiscountId.
 - If a provider applies a discount that cannot be represented by a FOCUS column, they SHOULD include additional columns to identify the source of the discount.
- Purchased discounts (e.g., *commitment discounts*) MUST be amortized.
 - The BilledCost MUST be 0 for any row where the commitment covers the entire cost for the charge period.
 - The EffectiveCost MUST include the portion of the amortized purchase cost that applies to this row.
 - The sum of the EffectiveCost for all rows where CommitmentDiscountStatus is "Used" or "Unused" for each CommitmentDiscountId over the entire duration of the commitment MUST be the same as the total BilledCost of the *commitment discount*.
 - The CommitmentDiscountId and ResourceId MUST be set to the ID assigned to the *commitment discount*. ChargeCategory MUST be set to "Purchase" on rows that represent a purchase of a *commitment discount*.
 - CommitmentDiscountStatus MUST be "Used" for ChargeCategory "Usage" rows that received a reduced price from a commitment. CommitmentDiscountId MUST be set to the ID assigned to the discount. ResourceId MUST be set to the ID of the resource that received the discount.
 - If a commitment is not fully utilized, the provider MUST include a row that represents the unused portion of the commitment for that *charge period*. These rows MUST be represented with CommitmentDiscountStatus set to "Unused" and ChargeCategory set to "Usage". Such rows MUST have their CommitmentDiscountId and ResourceId set to the ID assigned to the *commitment discount*.
- Credits that are applied after the fact MUST use a ChargeCategory of "Credit".

4.4.5. Exceptions

None

4.4.6. Introduced (version)

1.0-preview

4.5. Key-Value Format

Columns that provide Key-Value information are often used in place of separate columns for enumerating data which would be inherently sparse and/or without predetermined keys. This consolidates related information and provides more consistency in the schema. Key-value pairs are also referred to as name-value pairs, attribute-value pairs, or field-value pairs.

All key-value related columns defined in the FOCUS specification MUST follow the key-value formatting requirements listed below.

4.5.1. Attribute ID

KeyValueFormat

4.5.2. Attribute Name

Key-Value Format

4.5.3. Description

Rules and formatting requirements for columns appearing in a [FOCUS dataset](#) that convey data as key-value pairs.

4.5.4. Requirements

- Key-Value Format columns MUST contain a serialized JSON string, consistent with the [ECMA 404](#) definition of an object.
- Keys in a key-value pair MUST be unique within an object.
- Values in a key-value pair MUST be one of the following types: number, string, true, false, or null.
- Values in a key-value pair MUST NOT be an object or an array.

4.5.5. Exceptions

None

4.5.6. Introduced (version)

1.0-preview

4.6. Null Handling

Cost data [rows](#) that don't have a value that can be presented for a column must be handled in a consistent way to reduce friction for FinOps practitioners who consume the data for analysis, reporting, and other use cases.

All columns defined in the FOCUS specification MUST follow the null handling requirements listed below. Custom columns SHOULD also follow the same formatting requirements.

4.6.1. Attribute ID

NullHandling

4.6.2. Attribute Name

Null Handling

4.6.3. Description

Indicates how to handle columns that don't have a value.

4.6.4. Requirements

- Columns MUST use NULL when there isn't a value that can be specified for a nullable column.
- Columns MUST NOT use empty strings or placeholder values such as 0 for numeric columns or "Not Applicable" for string columns to represent a null or not having a value, regardless of whether the column allows nulls or not.

4.6.5. Exceptions

None

4.6.6. Introduced (version)

0.5

4.7. Numeric Format

Columns that provide numeric values conforming to specified rules and formatting requirements ensure clarity, accuracy, and ease of interpretation for humans and systems. The FOCUS specification does not require a specific level of precision for numeric values. The level of precision required for a given column is determined by the provider and should be part of a data definition published by the provider.

All columns capturing a numeric value, defined in the FOCUS specification, MUST follow the formatting requirements listed below. Custom numeric value capturing columns SHOULD adopt the same format requirements over time.

4.7.1. Attribute ID

NumericFormat

4.7.2. Attribute Name

Numeric Format

4.7.3. Description

Rules and formatting requirements for numeric columns appearing in a [FOCUS dataset](#).

4.7.4. Requirements

- Columns with a Numeric value format MUST contain a single numeric value.
- Numeric values MUST be expressed as an integer value, a decimal value, or a value expressed in scientific notation. Fractional notation MUST NOT be used.
- Numeric values expressed using scientific notation MUST be expressed using E notation "mEn" with a real number m and an integer n indicating a value of "m x 10^n". The sign of the exponent MUST only be expressed as part of the exponent value if n is negative.
- Numeric values MUST NOT be expressed with mathematical symbols, functions, or operators.
- Numeric values MUST NOT contain qualifiers or additional characters (e.g., currency symbols, units of measure, etc.).
- Numeric values MUST NOT contain commas or punctuation marks except for a single decimal point (".") if required to express a decimal value.
- Numeric values MUST NOT include a character to represent a sign for a positive value. A negative sign (-) MUST indicate a negative value.
- Columns with a Numeric value format MUST present one of the following values as the "Data type" in the column definition.

- Allowed values:

Data Type	Type Description
Integer	Specifies a numeric value represented by a whole number or by zero. Integer number formats correspond to standard data types defined by ISO/IEC 9899:2018
Decimal	Specifies a numeric value represented by a decimal number. Decimal formats correspond to ISO/IEC/IEEE 60559:2011 and IEEE 754-2008 definitions.

- Providers SHOULD define precision and scale for Numeric Format columns using one of the following precision values in a data definition document that providers publish.

- Allowed values:

Data Type	Precision	Definition	Range / Significant Digits
-----------	-----------	------------	----------------------------

Data Type	Precision	Definition	Range / Significant Digits
Integer	Short	16-bit signed short int ISO/IEC 9899:2018	-32,767 to +32,767
Integer	Long	32-bit signed long int ISO/IEC 9899:2018	-2,147,483,647 to +2,147,483,647
Integer	Extended	64-bit signed two's complement integer <i>or higher</i>	-(2 ⁶³ - 1) to (2 ⁶³ - 1)
Decimal	Single	32-bit binary format IEEE 754-2008 floating-point (decimal32)	9
Decimal	Double	64-bit binary format IEEE 754-2008 floating-point (decimal64)	16
Decimal	Extended	128-bit binary format IEEE 754-2008 floating-point (decimal128) or higher	36+

4.7.4.1. Examples

This format requires that single numeric values be represented using an integer or decimal format without additional characters or qualifiers. The following lists provide examples of values that meet the requirements and those that do not.

- Values Meeting Numeric Requirements:
 - -100.2
 - -3
 - 4
 - 35.2E-7
 - 1.234
- Values NOT Meeting Numeric Requirements
 - 1 1/2 - contains fractional notation
 - 35.2E+7 - contains a positive exponent with a sign
 - 35.24 x 10⁷ - contains an invalid format for scientific notation
 - [3,5,8] - contains an array
 - [4:5] - contains a range
 - 5i + 4 - contains a complex number
 - sqrt(2) - contains a mathematical symbol or operation
 - 2.3³ - contains an exponent
 - 32 GiB - contains a unit of measure
 - \$32 - contains a currency symbol
 - 3,432,342 - contains a comma
 - +333 - contains a positive sign

4.7.5. Exceptions

None

4.7.6. Introduced (version)

4.8. String Handling

Columns that capture string values conforming to specified requirements foster data integrity, interoperability, and consistency, improve data analysis and reporting, and support reliable data-driven decision-making.

All columns capturing a string value, defined in the FOCUS specification, MUST follow the requirements listed below. Custom string value capturing columns SHOULD adopt the same requirements over time.

4.8.1. Attribute ID

StringHandling

4.8.2. Attribute Name

String Handling

4.8.3. Description

Requirements for string-capturing columns appearing in a [FOCUS dataset](#).

4.8.4. Requirements

- String values MUST maintain the original casing, spacing, and other relevant consistency factors as specified by providers and end-users.
- [Charges](#) to mutable entities (e.g., resource names) MUST be accurately reflected in corresponding *charges* incurred after the change and MUST NOT alter *charges* incurred before the change, preserving data integrity and auditability for all *charge* records.
- Immutable string values that refer to the same entity (e.g., resource identifiers, region identifiers, etc.) MUST remain consistent and unchanged across all [billing periods](#).
- Empty strings and strings consisting solely of spaces SHOULD NOT be used in not-nullable string columns.

4.8.5. Exceptions

- When a record is provided after a change to a mutable string value and the [ChargeClass](#) is "Correction", the record MAY contain the altered value.

4.8.6. Introduced (version)

1.0

4.9. Unit Format

Billing data frequently captures data measured in units related to data size, count, time, and other [dimensions](#). The Unit Format attribute provides a standard for expressing units of measure in columns appearing in a [FOCUS dataset](#).

All columns defined in FOCUS specifying Unit Format as a value format MUST follow the requirements listed below.

4.9.1. Attribute ID

UnitFormat

4.9.2. Attribute Name

Unit Format

4.9.3. Description

Indicates standards for expressing measurement units in columns appearing in a *FOCUS dataset*.

4.9.4. Requirements

- Units SHOULD be expressed as a single unit of measure adhering to one of the following three formats.
 - <plural-units> - "GB", "Seconds"
 - <singular-unit>-<plural-time-units> - "GB-Hours", "MB-Days"
 - <plural-units>/<singular-time-unit> - "GB/Hour", "PB/Day"
- Units MAY be expressed with a unit quantity or time interval. If a unit quantity or time interval is used, the unit quantity or time interval MUST be expressed as a whole number. The following formats are valid:
 - <quantity> <plural-units> - "1000 Tokens", "1000 Characters"
 - <plural-units>/<interval> <plural-time-units> - "Units/3 Months"
- Unit values and components of columns using the Unit Format MUST use a capitalization scheme that is consistent with the capitalization scheme used in this attribute if that term is listed in this section. For example, a value of "gigabyte-seconds" would not be compliant with this specification as the terms "gigabyte" and "second" are listed in this section with the appropriate capitalization. If the unit is not listed in the table, it is to be used over a functional equivalent with a similar meaning with the same capitalization scheme.
- Units SHOULD be composed of the list of recommended units listed in this section unless the unit value covers a *dimension* not listed in the recommended unit set, or if the unit covers a count-based unit distinct from recommended values in the count *dimension* listed in this section.

4.9.4.1. Data Size Unit Names

Data size unit names MUST be abbreviated using one of the abbreviations in the following table. For example, a unit name of "TB" is a valid unit name, and a unit name of "terabyte" is an invalid unit name. Data size abbreviations can be considered both the singular and plural form of the unit. For

example, "GB" is both the singular and plural form of the unit "gigabyte", and "GBs" would be an invalid unit name. Values that exceed 10^{18} MUST use the abbreviation for exabit, exabyte, exbibit, and exbibyte, and values smaller than a byte MUST use the abbreviation for bit or byte. For example, the abbreviation "YB" for "yottabyte" is not a valid data size unit name as it represents a value larger than what is listed in the following table.

The following table lists the valid abbreviations for data size units from a single bit or byte to 10^{18} bits or bytes.

Data size in bits	Data size in bytes
b (bit) = 10^1	B (byte = 10^1)
Kb (kilobit = 10^3)	KB (kilobyte = 10^3)
Mb (megabit = 10^6)	MB (megabyte = 10^6)
Gb (gigabit = 10^9)	GB (gigabyte = 10^9)
Tb (terabit = 10^{12})	TB (terabyte = 10^{12})
Pb (petabit = 10^{15})	PB (petabyte = 10^{15})
Eb (exabit = 10^{18})	EB (exabyte = 10^{18})
Kib (kibibit = 2^{10})	KiB (kibibyte = 2^{10})
Mib (mebibit = 2^{20})	MiB (mebibyte = 2^{20})
Gib (gibibit = 2^{30})	GiB (gibibyte = 2^{30})
Tib (tebibit = 2^{40})	TiB (tebibyte = 2^{40})
Pib (pebibit = 2^{50})	PiB (pebibyte = 2^{50})
Eib (exbibit = 2^{60})	EiB (exbibyte = 2^{60})

4.9.4.2. Count-based Unit Names

A count-based unit is a noun that represents a discrete number of items, events, or actions. For example, a count-based unit can be used to represent the number of requests, instances, tokens, or connections.

If the following list of recommended values does not cover a count-based unit, a provider MAY introduce a new noun representing a count-based unit. All nouns appearing in units that are not listed in the recommended values table will be considered count-based units. A new count-based unit value MUST be capitalized.

Count
Count
Unit
Request
Token
Connection
Certificate
Domain
Core

4.9.4.3. Time-based Unit Names

A time-based unit is a noun that represents a time interval. Time-based units can be used to measure consumption over a time interval or in combination with another unit to capture a rate of consumption. Time-based units MUST match one of the values listed in the following table.

Time
Year
Month
Day
Hour
Minute
Second

4.9.4.4. Composite Units

If the unit value is a composite value made from combinations of one or more units, each component MUST also align with the set of recommended values.

Instead of "per" or "-" to denote a Composite Unit, slash ("/") and space(" ") MUST be used as a common convention. Count-based units like requests, instances, and tokens SHOULD be expressed using a value listed in the count *dimension*. For example, if a usage unit is measured as a rate of requests or instances over a period of time, the unit SHOULD be listed as "Requests/Day" to signify the number of requests per day.

4.9.5. Exceptions

None

4.9.6. Introduced (version)

1.0-preview

5. Metadata

The FOCUS specification defines a metadata structure to be supplied by data providers to facilitate practitioners' use of FOCUS data. This metadata includes general information about the data generator and the schema of the [FOCUS dataset](#).

FOCUS Metadata SHOULD be provided in a format that is accessible programmatically, such as a file, website, API, or table. Providers SHOULD provide documentation on their implementation of the FOCUS metadata.

5.1. Data Generator

The FOCUS metadata about the generator of the FOCUS data.

5.1.1. Requirements

The FOCUS Data Generator metadata MUST be provided. This metadata MUST be of type Object and MUST NOT contain null values.

5.1.2. Schema Example

For an example of the FOCUS Data Generator metadata please refer to: [Data Generator Example](#)

5.1.3. Data Generator

Human-readable name of the entity that is generating the data.

The DataGenerator MUST be provided in the metadata. DataGenerator MUST be of type String and MUST NOT be null. The DataGenerator SHOULD be easily associated with the provider who generated the [FOCUS dataset](#).

5.1.3.1. Metadata ID

DataGenerator

5.1.3.2. Metadata Name

Data Generator

5.1.3.3. Content constraints

Constraint	Value
------------	-------

Constraint	Value
Feature level	Mandatory
Allows nulls	False
Data type	String
Value format	<not specified>

5.1.3.4. Introduced (version)

1.0

5.2. Schema

The schema metadata object and its contents provides information about the structure of the data provided.

5.2.1. Requirements

5.2.1.1. Reference to FOCUS Data

FOCUS data artifacts, whether they are data files, data streams, or data tables, **MUST** provide a clear reference to the schema of the data. This reference **MUST** be retrievable without inspection of the contents of the FOCUS data within the data artifact. For some delivery mechanisms such as database tables, the provider may rely on the schema functionality of the providing system.

It is recommended that the schema reference be provided as an external reference rather than included in full as metadata accompanying the data artifact. This allows for easier understanding of when changes to the schema of the [FOCUS datasets](#) occurs.

5.2.1.2. Schema Metadata Creation

Should the provider change the structure of the supplied FOCUS data artifact, a new schema metadata object **MUST** be supplied. These scenarios include, but are not limited to:

- [Adding a new column](#)
- [Removing a column](#)
- [Changing column metadata](#)
- [FOCUS Version has changed](#)
- [Data Generator Version has changed](#)
- [Correcting schema metadata errors](#)

5.2.1.3. Schema Metadata Updates

Should there be an error where the schema metadata object does not match the schema of the FOCUS data artifact, the provider **MUST** update the schema metadata object to match the schema of the FOCUS data artifact. This is to ensure that the schema metadata object is always accurate.

5.2.2. Schema Example

For an example of the FOCUS schema metadata please refer to: [Schema Metadata Example](#)

5.2.3. Schema ID

The Schema ID provides the reference item to associate which Schema was used for the generation of a FOCUS Dataset.

The Schemald MUST be present in the metadata. The Schemald MUST be of String. It is RECOMMENDED for Schemald to be a Globally Unique Identifier (GUID).

5.2.3.1. Metadata ID

Schemald

5.2.3.2. Metadata Name

Schema ID

5.2.3.3. Content constraints

Constraint	Value
Feature level	Mandatory
Allows nulls	False
Data type	STRING
Value format	Recommend GUID String

5.2.3.4. Introduced (version)

1.0

5.2.4. Creation Date

Date the schema was created.

The CreationDate MUST be present in the metadata. This MUST be of type Date/Time and MUST NOT contain null values. CreationDate MUST conform to [DateTimeFormat](#).

5.2.4.1. Metadata ID

CreationDate

5.2.4.2. Metadata Name

Creation Date

5.2.4.3. Content constraints

Constraint	Value
Feature level	Mandatory
Allows nulls	False
Data type	Date/Time
Value format	Date/Time Format

5.2.4.4. Introduced (version)

1.0

5.2.5. FOCUS Version

The version of FOCUS utilized for building the dataset.

The FocusVersion MUST be provided in the metadata. FocusVersion MUST be of type String and MUST NOT contain null values. FocusVersion MUST match one of the published versions of the FOCUS specification. FocusVersion MUST match the version of the FOCUS specification that the [FOCUS dataset](#) conforms to.

5.2.5.1. Metadata ID

FocusVersion

5.2.5.2. Metadata Name

FOCUS Version

5.2.5.3. Content constraints

Constraint	Value
Feature level	Mandatory
Allows nulls	False
Data type	STRING
Value format	Must align with a published FocusVersion

5.2.5.4. Introduced (version)

5.2.6. Data Generator Version

The DataGeneratorVersion MAY be supplied to declare the version of logic by which the [FOCUS dataset](#) was generated and is separate from FOCUS Version. DataGeneratorVersion allows for the provider to specify changes that may not result in a structural change in the data. It is suggested that the DataGeneratorVersion use a versioning approach such as [SemVer](#) version.

The DataGeneratorVersion column adheres to the following requirements:

- DataGeneratorVersion MAY be present in FOCUS metadata.
- DataGeneratorVersion MUST be of type String.
- DataGeneratorVersion MUST conform to [StringHandling](#) requirements.
- DataGeneratorVersion MUST NOT be null.
- When FocusVersion is changed, a new DataGeneratorVersion MUST be also changed.
- Data generators MUST document what changes are present in the DataGeneratorVersion.

5.2.6.1. Metadata ID

DataGeneratorVersion

5.2.6.2. Metadata Name

Data Generator Version

5.2.6.3. Content constraints

Constraint	Value
Feature level	Optional
Allows nulls	False
Data type	STRING
Value format	<not specified>

5.2.6.4. Introduced (version)

1.1

5.2.7. Column Definition

The FOCUS metadata schema column definition provides a list of the columns present in the [FOCUS dataset](#) along with metadata about the columns.

5.2.7.1. Requirements

This metadata MUST be present in the FOCUS metadata schema. This metadata MUST be of type Object and MUST NOT contain null values.

5.2.7.2. Column Name

The name of the column provided in the [FOCUS dataset](#).

The ColumnName MUST be provided in the FOCUS Metadata schema. ColumnName MUST be of type String and MUST NOT contain null values.

5.2.7.2.1. Metadata ID

ColumnName

5.2.7.2.2. Metadata Name

Column Name

5.2.7.2.3. Content constraints

Constraint	Value
Feature level	Mandatory
Allows nulls	False
Data type	String
Value format	<not specified>

5.2.7.2.4. Introduced (version)

1.0

5.2.7.3. Data Type

The data type of the column provided in the [FOCUS dataset](#).

The DataType MUST be provided in the FOCUS Metadata schema. DataType MUST be of type String and MUST NOT contain null values.

5.2.7.3.1. Metadata ID

DataType

5.2.7.3.2. Metadata Name

Data Type

5.2.7.3.3. Content constraints

Constraint	Value
Feature level	Mandatory
Allows nulls	False
Data type	String
Value format	<not specified>

5.2.7.3.4. Introduced (version)

1.0

5.2.7.4. Deprecated

The deprecation status of any column in a [FOCUS dataset](#).

Deprecated MUST be provided in the FOCUS Metadata schema when a column will be removed in a future delivered schema definition. DataType MUST be of type Boolean and MUST NOT contain null values. The value of deprecated should only be "true" if the column is deprecated. Providers can choose to always provide the deprecation key or elect to only include it when the deprecation status of a column is "true". Deprecated must be "true" when the provider removes a column at a future date, or the column has been identified for deprecation for the FOCUS version identified in the schema definition.

5.2.7.4.1. Metadata ID

Deprecated

5.2.7.4.2. Metadata Name

Deprecated

5.2.7.4.3. Content constraints

Constraint	Value
Feature level	Conditional
Allows nulls	False
Data type	Boolean
Value format	<not specified>

5.2.7.4.4. Introduced (version)

1.2

5.2.7.5. Numeric Precision

Numeric Precision is the maximum number of digits for the values in the column.

NumericPrecision SHOULD be provided in the FOCUS Metadata schema for Numeric Format columns.

NumericPrecision MUST be of type Integer and MUST NOT contain null values.

5.2.7.5.1. Metadata ID

NumericPrecision

5.2.7.5.2. Metadata Name

Numeric Precision

5.2.7.5.3. Content constraints

Constraint	Value
Feature level	Conditional
Allows nulls	False
Data type	Integer
Value format	Numeric Format

5.2.7.5.4. Introduced (version)

1.0

5.2.7.6. Number Scale

The number scale of the data provides the maximum number of digits after the decimal point in decimal numbers.

NumberScale SHOULD be provided in the FOCUS Metadata schema for Decimal columns.

NumberScale MUST be of type Integer and MUST NOT contain null values.

5.2.7.6.1. Metadata ID

NumberScale

5.2.7.6.2. Metadata Name

Number Scale

5.2.7.6.3. Content constraints

Constraint	Value
Feature level	Conditional
Allows nulls	False
Data type	Integer
Value format	Numeric Format

5.2.7.6.4. Introduced (version)

1.0

5.2.7.7. PreviousColumnName

The PreviousColumnName field indicates that on that schema the column where the key is included was renamed.

In cases where the PreviousColumnName is present, the following applies:

- PreviousColumnName MUST not be null.
- PreviousColumnName MUST be of type String.
- PreviousColumnName MUST be the name used in previous versions of the schema.
- PreviousColumnName MUST NOT be present in schema versions created after the rename.

5.2.7.7.1. Metadata ID

PreviousColumnName

5.2.7.7.2. Metadata Name

Previous Column Name

5.2.7.7.3. Content constraints

Constraint	Value
Feature level	Conditional
Allows nulls	False
Data type	String
Value format	<not specified>

5.2.7.7.4. Introduced (version)

1.2

5.2.7.8. Provider Tag Prefixes

The Provider Tag Prefixes define the list of prefixes used in the tag name of provider-defined [tags](#). This metadata is useful for the consumer to identify which tags are provider-defined vs user-defined.

The ProviderTagPrefixes MUST be provided when ColumnName is equal to Tags. The ProviderTagPrefix MUST be of type Array of Strings. The ProviderTagPrefixes SHOULD be easily associated with the provider who generated the [FOCUS dataset](#).

5.2.7.8.1. Metadata ID

ProviderTagPrefixes

5.2.7.8.2. Metadata Name

Provider Tag Prefixes

5.2.7.8.3. Content constraints

Constraint	Value
Feature level	Conditional
Allows nulls	False
Data type	Array
Value format	STRING datatype values in the array

5.2.7.8.4. Introduced (version)

1.0

5.2.7.9. String Encoding

The string encoding scheme of the column provided in the [FOCUS dataset](#).

StringEncoding SHOULD be provided in the FOCUS Metadata schema when it is required to know this information in order to successfully read the data. StringEncoding MUST be of type String and MUST NOT contain null values.

5.2.7.9.1. Metadata ID

StringEncoding

5.2.7.9.2. Metadata Name

StringEncoding

5.2.7.9.3. Content constraints

Constraint	Value
Feature level	Conditional
Allows nulls	False
Data type	String
Value format	<not specified>

5.2.7.9.4. Introduced (version)

1.0

5.2.7.10. String Max Length

The string max length of the data that can be stored in the column.

StringMaxLength SHOULD be provided in the FOCUS Metadata schema for String columns.
StringMaxLength MUST be of type Integer and MUST NOT contain null values.

5.2.7.10.1. Metadata ID

StringMaxLength

5.2.7.10.2. Metadata Name

String Max Length

5.2.7.10.3. Content constraints

Constraint	Value
Feature level	Conditional
Allows nulls	False
Data type	Integer
Value format	Numeric Format

5.2.7.10.4. Introduced (version)

1.0

6. Use Case Library

This specification is based on a set of common FinOps use cases, which are publicly available at <https://focus.finops.org/use-cases/>. Developed by FinOps practitioners, these use cases are organized by persona and capability, making it easy to find relevant scenarios. Each use case includes sample SQL queries to help you get started with implementation.

7. Glossary

Adjustment

A charge representing a modification to billing data to account for certain events or circumstances not previously captured, or captured incorrectly. Examples include billing errors, service disruptions, or pricing changes.

Amortization

The distribution of upfront costs over time to accurately reflect the consumption or benefit derived from the associated resources or services. Amortization is valuable when the commitment period (time duration of the cost) extends beyond the granularity of the source report.

Availability Zone

A collection of geographically separated locations containing a data center or cluster of data centers. Each availability zone (AZ) should have its own power, cooling, and networking, to provide redundancy and fault tolerance.

Billed Cost

A charge that serves as the basis for invoicing. It includes the total amount of fees and discounts, signifying a monetary obligation. Valuable when reconciling cash outlay with incurred expenses is required, such as cost allocation, budgeting, and invoice reconciliation.

Billing Account

A container for resources and/or services that are billed together in an invoice. A billing account may have sub accounts, all of whose costs are consolidated and invoiced to the billing account.

Billing Currency

An identifier that represents the currency that a charge for resources and/or services was billed in.

Billing Period

The time window that an organization receives an invoice for, inclusive of the start date and exclusive of the end date. It is independent of the time of usage and consumption of resources and services.

Block Pricing

A pricing approach where the cost of a particular resource or service is determined based on predefined quantities or tiers of usage. In these scenarios, the Pricing Unit and the corresponding Pricing Quantity can be different from the Consumed Unit and Consumed Quantity.

Capacity Reservation

A capacity reservation is an agreement that secures a dedicated amount of resources or services for a specified period. This ensures the reserved capacity is always available and accessible, even if it's not fully utilized. Customers are typically charged for the reserved capacity, regardless of actual consumption.

Charge

A row in a FOCUS-compatible cost and usage dataset.

Charge Period

The time window for which a charge is effective, inclusive of the start date and exclusive of the end date. The charge period for continuous usage should match the time granularity of the dataset (e.g., 1 hour for hourly, 1 day for daily). The charge period for a non-usage charge with time boundaries should match the duration of eligibility.

Cloud Service Provider (CSP)

A company or organization that provides remote access to computing resources, infrastructure, or applications for a fee.

Commitment

A customer's agreement to consume a specific quantity of a service or resource over a defined period, usually also creating a financial commitment throughout the entirety of the commitment period. Some commitments also hold Providers to certain assurance levels of resource availability.

Commitment Discount

A billing discount model that offers reduced rates on preselected SKUs in exchange for an obligated usage or spend amount over a predefined term. Commitment discount purchases, made upfront and/or with recurring monthly payments are amortized evenly across predefined charge periods (i.e., hourly), and unused amounts cannot be carried over to subsequent charge periods. Commitment discounts are publicly available to customers without special contract arrangements.

Commitment Discount Flexibility

A feature of [commitment discounts](#) that may further transform the predetermined amount of usage purchased or consumed based on additional, provider-specific requirements.

Contracted Unit Price

The agreed-upon unit price for a single [Pricing Unit](#) of the associated SKU, inclusive of negotiated discounts, if present, and exclusive of any other discounts. This price is denominated in the [Billing Currency](#).

Correction

A charge to correct cost or usage data in a previously invoiced [billing period](#).

Credit

A financial incentive or allowance granted by a provider unrelated to other past/current/future charges.

Dimension

A specification-defined categorical attribute that provides context or categorization to billing data.

Effective Cost

The amortized cost of the charge after applying all reduced rates, discounts, and the applicable portion of relevant, prepaid purchases (one-time or recurring) that covered this charge.

Exclusive End Bound

A Date/Time Format value that is not contained within the ending bound of a time period.

Finalized Tag

A tag with one tag value chosen from a set of possible tag values after being processed by a set of provider-defined or user-defined rules.

FinOps Cost and Usage Specification (FOCUS)

An open-source specification that defines requirements for billing data.

FOCUS Dataset

A structured collection of cost and usage data that meets the [BCP14](#) criteria defined by FOCUS. In addition to FOCUS columns, the dataset should include custom provider columns (prefixed with x_) when these columns provide additional information not captured by the existing FOCUS columns. If introducing a custom column could result in splitting original charge records into multiple entries, the data generator is responsible for ensuring that the FOCUS dataset fully conforms to all aggregation-related requirements for metric columns, particularly those concerning costs and quantities.

Inclusive Start Bound

A Date/Time Format value that is contained within the beginning bound of a time period.

Interruptible

A category of compute resources that can be paused or terminated by the CSP within certain criteria, often advertised at reduced unit pricing when compared to the equivalent non-interruptible resource.

List Unit Price

The suggested provider-published unit price for a single [Pricing Unit](#) of the associated [SKU](#), exclusive of any discounts. This price is denominated in the [Billing Currency](#).

Managed Service Provider (MSP)

A company or organization that provides outsourced management and support of a range of IT services, such as network infrastructure, cybersecurity, cloud computing, and more.

Metric

A FOCUS-defined column that provides numeric values, allowing for aggregation operations such as arithmetic operations (sum, multiplication, averaging etc.) and statistical operations.

National Currency

A government-issued currency (e.g., US dollars, Euros).

Negotiated Discount

A contractual agreement where a customer commits to specific spend or usage goals over a [term](#) in exchange for discounted rates across varying SKUs. Unlike [commitment discounts](#), negotiated discounts are typically more customized to customer's accounts, can be utilized at varying frequencies, and may overlap with [commitment discounts](#).

On-Demand

A term that describes a service that is available and provided immediately or as needed, without requiring a pre-scheduled appointment or prior arrangement. In cloud computing, virtual machines can be created and terminated as needed, i.e., on demand.

Pascal Case

Pascal Case (PascalCase, also known as UpperCamelCase) is a format for identifiers which contain one or more words meaning the words are concatenated together with no delimiter and the first letter of each word is capitalized.

Potato

A long and often painful conversation had by the FOCUS contributors. Sometimes the name of a thing that we could not yet name. No starchy root vegetables were harmed during the production of this specification. We thank potato for its contribution in the creation of this specification.

Practitioner

An individual who performs FinOps within an organization to maximize the business value of using cloud and cloud-like services.

Price List

A comprehensive list of prices offered by a provider.

Provider

An entity that made internal or 3rd party resources and/or services available for purchase.

Refund

A return of funds that have previously been charged.

Resource

A unique component that incurs a charge.

Row

A row in a FOCUS-compatible cost and usage dataset.

Service

An offering that can be purchased from a provider, and can include many types of usage or other charges; eg., a cloud database service may include compute, storage, and networking charges.

SKU

A construct composed of the common properties of a product offering associated with one or many SKU Prices.

SKU Price

A pricing construct that encompasses SKU properties (e.g., functionality and technical specifications), along with core stable pricing details for a particular SKU, while excluding dynamic or negotiable pricing elements such as unit price amounts, currency (and related exchange rates), temporal validity (e.g., effective dates), and contract- or negotiation-specific factors (e.g., contract or account identifiers, and negotiable discounts).

Sub Account

A sub account is an optional provider-supported construct for organizing resources and/or services connected to a billing account. Sub accounts must be associated with a billing account as they do not receive invoices.

Tag

A metadata label assigned to a resource to provide information about it or to categorize it for organizational and management purposes.

Tag Source

A Resource or Provider-defined construct for grouping resources and/or other Provider-defined construct that a Tag can be assigned to.

Term

A duration of a contractual agreement like with a [commitment discount](#) or [negotiated discount](#).

Virtual Currency

A proprietary currency (e.g., credits, tokens) issued by providers and independent of government regulation.

8. Appendix

This section is non-normative.

8.1. Examples: Commitment Discount Scenarios

A [commitment discount](#) is a billing discount model that offers reduced rates on preselected [SKUs](#) in exchange for an obligated usage or spend amount over a predefined [term](#). *Commitment discounts* typically consist of purchase and usage records within cost and usage datasets.

Usage-based *commitment discounts* obligate a customer to a predetermined amount of usage over a preselected *term*. In some cases, usage-based *commitment discounts* also feature [commitment discount flexibility](#) which may expand the types of [resources](#) that a *commitment discount* can cover. It is important to note when mixing *commitment discounts* with and without *commitment discount flexibility*, the [CommitmentDiscountUnit](#) should reflect this difference.

Spend-based commitment discounts obligate a customer to a predetermined amount of spend over a preselected *term*. In the usage examples below, each [row](#) measures the monetary amount of the hourly commit consumed by the *commitment discount*, so the [CommitmentDiscountUnit](#) chosen is "USD", or the [billing currency](#).

8.1.1. Purchasing

While customers are bound to the *term* of a *commitment discount*, providers offer some or all of the following payment options before and/or during the *term*:

- *All Upfront* - The *commitment discount* is paid in full before the *term* begins.
- *No Upfront* - The *commitment discount* is paid on a repeated basis, typically over each [billing period](#) of the *term*.
- *Partial Upfront* - Some of the *commitment discount* is paid before the *term* begins, and the rest is paid repeatedly over the *term*.

For example, if a customer buys a 1-year, spend-based *commitment discount* with a \$1.00 hourly commit and pays with the partial option, the *commitment discount's* payment consists of a one-time purchase in the beginning of the *term* and monthly recurring purchases with the following totals:

1. *One-Time* - \$4,380 (24 hours * 365 days * $1.00 * 0.5)
2. *Recurring* - \$182.50 (24 hours * 365 days * $1.00 / 12 months)

8.1.2. Usage

Commitment discounts follow a "use-it-or-lose-it" model where the [amortization](#) of a *commitment discount's* purchase applies evenly to eligible *resources* over each [charge period](#) of the *term*.

For example, if a customer buys a spend-based *commitment discount* with a \$1.00 hourly commit in January (31 days), only \$1.00 is eligible for consumption for each hourly *charge period*. If a customer has eligible *resources* running during this *charge period*, an amount of up to \$1.00 will be allocated to these *resources*. Conversely, if a customer does not have eligible *resources* running that fully take advantage of this \$1.00 during this *charge period*, then some or all of this amount will go to waste.

8.1.3. Commitment Discounts in FOCUS

Within the FOCUS specification, the following examples demonstrate how a *commitment discount* appears across various payment and usage scenarios.

8.1.3.1. Purchase Rows

All *commitment discount* purchases appear with a positive [BilledCost](#), [PricingCategory](#) as "Standard", and with the *commitment discount's* id populating both the [Resourcelid](#) and [CommitmentDiscountId](#) value. One-time purchases appear as a single record with [ChargeCategory](#) as "Purchase", [ChargeFrequency](#) as "One-Time", and the total quantity and units for *commitment discount's term* reflected as [CommitmentDiscountQuantity](#) and [CommitmentDiscountUnit](#), respectively.

Recurring purchases are allocated across all corresponding *charge periods* of the *term* when [ChargeCategory](#) is "Purchase", [ChargeFrequency](#) is "Recurring", and [CommitmentDiscountQuantity](#) and [CommitmentDiscountUnit](#) are reflected only for that *charge period*.

Using the same *commitment discount* example as above with a one-year, spend-based *commitment discount* with a \$1.00 hourly commit purchased on Jan 1, 2023, various purchase options are available:

8.1.3.1.1. Scenario #1: All Upfront

The entire *commitment discount* is billed *once* during the first *charge period* of the *term* for \$8,670 (derived as 24 hours * 365 days * $1.00).

[CSV Example](#)

8.1.3.1.2. Scenario #2: No Upfront

The *commitment discount* is billed across all 8,760 (24 hours * 365 days) *charge periods* of the *term* with \$1.00 allocated to each *charge period* over the *term*.

[CSV Example](#)

This example shows the first three hourly rows of 8,760 total rows that are all the same except for the incrementing monthly and hourly timeframes denoted in the Billing Period and Charge Period columns, respectively.

8.1.3.1.3. Scenario #3: Partial Upfront

With a 50/50 split, half of the commitment is billed *once* during the first *charge period* of the *term* for \$4,380 (derived as 24 hours * 182.5 days * $1.00), and the other half is billed across each *charge period* over the term, derived as ($1.00 * 8,760 hours * 0.5). Amortized costs incur half of the amount (i.e., \$0.50) from the one-time purchase and the other half from the recurring purchase.

[CSV Example](#)

This example shows the first three hourly rows of 8,760 total rows that are all the same except for the incrementing monthly and hourly timeframes denoted in the Billing Period and Charge Period columns, respectively.

8.1.3.2. Usage Rows

Amortization of commitment discounts occur similarly regardless of how *commitment discount* purchases are made. The same usage-based or spend-based amount is applied evenly across all *charge periods* and potentially allocated to eligible *resources*. Continuing with the same *commitment discount* example, a one-year, spend-based *commitment discount* with a \$1.00 hourly commit and 1 *resource* (for simplicity) yields 4 types of scenarios that can occur during a *charge period*:

- Scenario #1: An eligible *resource* fully consumes the allocated amount (100% utilization)
- Scenario #2: No eligible *resource* consumes the allocated amount (0% utilization)
- Scenario #3: An eligible *resource* partially consumes the allocated amount (75% utilization)
- Scenario #4: An eligible *resource* fully consumes the \$1.00 hourly commit with an overage (100% utilization + overage)

8.1.3.2.1. Scenario #1: An eligible *resource* fully consumes the allocated amount (100% utilization)

In this scenario, one eligible *resource* runs for the full hour and consumes \$1.00, so one *row* allocated to the *resource* is produced.

[CSV Example](#)

8.1.3.2.2. Scenario #2: No eligible *resource* consumes the allocated amount (0% utilization)

In this situation, the full eligible, \$1.00 amount remained unutilized and results in 1 unused *row*. In this scenario, it is important to note that while *CommitmentDiscountQuantity* is not because \$1 was still drawn down by the *commitment discount* even though, no *resource* was allocated, so [ConsumedQuantity](#) and [ConsumedUnit](#) are null.

[CSV Example](#)

8.1.3.2.3. Scenario #3: An eligible *resource* partially consumes the allocated amount (75% utilization)

In this scenario, one eligible *resource* runs for the full hour and consumes \$0.75 of the \$1.00 allocation. One *row* shows \$0.75 to a *resource*, and the other *row* shows that \$0.25 was unused.

[CSV Example](#)

8.1.3.2.4. Scenario #4: An eligible *resource* fully consumes the \$1.00 hourly commit with an overage (100% utilization + overage)

In this scenario, one eligible *resource* runs for the full hour and is charged \$1.50. One *row* shows that \$1.00 was *amortized* from the *commitment discount*, and the other shows that \$0.50 was charged as standard, on-demand spend.

[CSV Example](#)

8.2. Examples: Commitment Discount Flexibility

A usage-based [commitment discount](#) obligates a customer to a usage amount for one or more related SKUs in return for reduced rates. For example, when a usage-based *commitment discount* is

purchased to cover a specific database SKU, this commitment will cover every hour over the term where at least one instance of this SKU is running. The usage-based commitment can cover 1 resource over the hour, or in the case of [commitment discount flexibility](#), it can cover a portion of 1 resource or multiple resources at a time.

When mixing usage-based commitment discounts with and without *commitment discount flexibility* and [CommitmentDiscountQuantity](#) measured by time, it is important to differentiate the [CommitmentDiscountUnit](#) for each type of *commitment discount*. In each scenario below, *commitment discounts without commitment discount flexibility* applied use "Hour" as the *CommitmentDiscountUnit*, and conversely commitment discounts *with commitment discount flexibility* applied use "Normalized Hour" as the *CommitmentDiscountUnit*.

For more details on exactly how *commitment discounts* purchase and usage rows appear with and without *commitment discount flexibility*, see the following scenarios:

8.2.1. 100% utilization without commitment discount flexibility

8.2.1.1. Context

For this example, fictitious provider, *TinyCloud*, offers the following SKU catalog which is used in the scenario below.

8.2.1.2. SKU Catalog

Service	Sku Id	Sku Price Id	Sku Price Unit Price	Normalization Factor
Compute	VM_SMALL	VM_SMALL_COMMITTED_PURCHASE_NO_UPFRONT	\$0.50	1
Compute	VM_MEDIUM	VM_MEDIUM_COMMITTED_PURCHASE_NO_UPFRONT	\$1.00	2
Compute	VM_LARGE	VM_LARGE_COMMITTED_PURCHASE_NO_UPFRONT	\$1.50	3
Compute	VM_XLARGE	VM_XLARGE_COMMITTED_PURCHASE_NO_UPFRONT	\$2.00	4
Compute	VM_SMALL	VM_SMALL_COMMITTED_HOUR	\$0.50	1
Compute	VM_MEDIUM	VM_MEDIUM_COMMITTED_HOUR	\$1.00	2
Compute	VM_LARGE	VM_LARGE_COMMITTED_HOUR	\$1.50	3
Compute	VM_XLARGE	VM_XLARGE_COMMITTED_HOUR	\$2.00	4
Compute	VM_SMALL	VM_SMALL_ON_DEMAND_HOUR	\$1.00	1
Compute	VM_MEDIUM	VM_MEDIUM_ON_DEMAND_HOUR	\$2.00	2
Compute	VM_LARGE	VM_LARGE_ON_DEMAND_HOUR	\$3.00	3
Compute	VM_XLARGE	VM_XLARGE_ON_DEMAND_HOUR	\$4.00	4

The above SKU Catalog shows that this provider only has 1 service that offers 4 virtual machine SKUs at various list rates, [commitment discount](#) rates, and normalization factors. Each SKU's normalization factor classifies its relative size to its *commitment discount* rate. Usage-based *commitment discounts* with [commitment discount flexibility](#) can fully cover any combination of 1 or more SKUs where the sum of their normalization factor is less than or equal to the normalization factor of the *commitment discount*.

8.2.1.3. Scenario

- 1 no upfront *commitment discount* is purchased for 1 year (2023) for 1 VM_LARGE.
- 1 VM_LARGE resource runs for 1 hour from 2023-01-01T00:00:00 to 2023-01-01T01:00:00.

8.2.1.4. Outcome

- 1 recurring, purchase record exists for 1 eligible "Hour" of the no upfront, *commitment discount* and incurs a \$1.50 [BilledCost](#).
- The *commitment discount* covers the first [charge period](#) for 1 VM_LARGE resource incurring a \$1.50 [EffectiveCost](#).

[CSV Example](#)

8.2.2. 0% utilization without commitment discount flexibility

8.2.2.1. Context

For this example, fictitious provider, *TinyCloud*, offers the following SKU catalog which is used in the scenario below.

8.2.2.2. SKU Catalog

Service	Sku Id	Sku Price Id	Sku Price Unit Price	Normalization Factor
Compute	VM_SMALL	VM_SMALL_COMMITTED_PURCHASE_NO_UPFRONT	\$0.50	1
Compute	VM_MEDIUM	VM_MEDIUM_COMMITTED_PURCHASE_NO_UPFRONT	\$1.00	2
Compute	VM_LARGE	VM_LARGE_COMMITTED_PURCHASE_NO_UPFRONT	\$1.50	3
Compute	VM_XLARGE	VM_XLARGE_COMMITTED_PURCHASE_NO_UPFRONT	\$2.00	4
Compute	VM_SMALL	VM_SMALL_COMMITTED_HOUR	\$0.50	1
Compute	VM_MEDIUM	VM_MEDIUM_COMMITTED_HOUR	\$1.00	2
Compute	VM_LARGE	VM_LARGE_COMMITTED_HOUR	\$1.50	3
Compute	VM_XLARGE	VM_XLARGE_COMMITTED_HOUR	\$2.00	4
Compute	VM_SMALL	VM_SMALL_ON_DEMAND_HOUR	\$1.00	1
Compute	VM_MEDIUM	VM_MEDIUM_ON_DEMAND_HOUR	\$2.00	2
Compute	VM_LARGE	VM_LARGE_ON_DEMAND_HOUR	\$3.00	3
Compute	VM_XLARGE	VM_XLARGE_ON_DEMAND_HOUR	\$4.00	4

The above SKU Catalog shows that this provider only has 1 service that offers 4 virtual machine SKUs at various list unit prices, [commitment discount](#) unit prices, and normalization factors. Each SKU's normalization factor classifies its relative size to its *commitment discount* unit price. Usage-based *commitment discounts* with [commitment discount flexibility](#) can fully cover any combination of 1 or more SKUs where the sum of their normalization factor is less than or equal to the normalization factor of the *commitment discount*.

8.2.2.3. Scenario

- 1 no upfront *commitment discount* is purchased for 1 year (2023) for 1 VM_LARGE.
- 1 VM_MEDIUM resource runs for 1 hour from 2023-01-01T00:00:00 to 2023-01-01T01:00:00.

8.2.2.4. Outcome

- 1 recurring, purchase record exists for 1 eligible "Hour" of the no upfront, *commitment discount* and incurs a \$1.50 [BilledCost](#).
- The VM_LARGE *commitment discount* is unused for the corresponding [charge period](#) because no VM_LARGE resources are running and incurs a \$1.50 [EffectiveCost](#).
- 1 hour of on-demand usage is incurred by the VM_MEDIUM resource and incurs a \$2.00 [BilledCost](#) and [EffectiveCost](#).

[CSV Example](#)

8.2.3. 100% utilization with commitment discount flexibility with 1 resource

8.2.3.1. Context

For this example, fictitious provider, *TinyCloud*, offers the following SKU catalog which is used in the scenario below.

8.2.3.2. SKU Catalog

Service	Sku Id	Sku Price Id	Sku Price Unit Price	Normalization Factor
Compute	VM_SMALL	VM_SMALL_COMMITTED_PURCHASE_NO_UPFRONT	\$0.50	1
Compute	VM_MEDIUM	VM_MEDIUM_COMMITTED_PURCHASE_NO_UPFRONT	\$1.00	2
Compute	VM_LARGE	VM_LARGE_COMMITTED_PURCHASE_NO_UPFRONT	\$1.50	3
Compute	VM_XLARGE	VM_XLARGE_COMMITTED_PURCHASE_NO_UPFRONT	\$2.00	4
Compute	VM_SMALL	VM_SMALL_COMMITTED_HOUR	\$0.50	1
Compute	VM_MEDIUM	VM_MEDIUM_COMMITTED_HOUR	\$1.00	2
Compute	VM_LARGE	VM_LARGE_COMMITTED_HOUR	\$1.50	3
Compute	VM_XLARGE	VM_XLARGE_COMMITTED_HOUR	\$2.00	4
Compute	VM_SMALL	VM_SMALL_ON_DEMAND_HOUR	\$1.00	1
Compute	VM_MEDIUM	VM_MEDIUM_ON_DEMAND_HOUR	\$2.00	2
Compute	VM_LARGE	VM_LARGE_ON_DEMAND_HOUR	\$3.00	3
Compute	VM_XLARGE	VM_XLARGE_ON_DEMAND_HOUR	\$4.00	4

The above SKU catalog shows that this provider only has 1 service that offers 4 virtual machine SKUs at various list rates, *commitment discount* rates, and normalization factors. Each SKU's normalization factor classifies its relative size to its *commitment discount* rate. Usage-based [commitment discounts](#) with [commitment discount flexibility](#) can fully cover any combination of 1 or more SKUs where the sum of their normalization factor is less than or equal to the normalization factor of the *commitment discount*.

8.2.3.3. Scenario

- 1 no upfront *commitment discount* is purchased for 1 year (2023) for 1 VM_SMALL which has a normalization factor of 1.
- 1 VM_LARGE resource runs for 1 hour from 2023-01-01T00:00:00 to 2023-01-01T01:00:00 with a normalization factor of 4.

8.2.3.4. Outcome

- 1 recurring, purchase record exists for 1 eligible "Normalized Hour" of the no upfront, *commitment discount* and incurs a \$0.50 [BilledCost](#).
- The VM_SMALL *commitment discount* is fully utilized within the corresponding [charge period](#), covers 25% of the VM_LARGE resource, and incurs a \$0.50 [EffectiveCost](#).
- The VM_LARGE resource incurs an additional, on-demand \$2.25 *BilledCost* and *EffectiveCost*.

[CSV Example](#)

8.2.4. 100% utilization with commitment discount flexibility with 2 resources

8.2.4.1. Context

For this example, fictitious provider, *TinyCloud*, offers the following SKU catalog which is used in the scenario below.

8.2.4.2. SKU Catalog

Service	Sku Id	Sku Price Id	Sku Price Unit Price	Normalization Factor
Compute	VM_SMALL	VM_SMALL_COMMITTED_PURCHASE_NO_UPFRONT	\$0.50	1
Compute	VM_MEDIUM	VM_MEDIUM_COMMITTED_PURCHASE_NO_UPFRONT	\$1.00	2
Compute	VM_LARGE	VM_LARGE_COMMITTED_PURCHASE_NO_UPFRONT	\$1.50	3
Compute	VM_XLARGE	VM_XLARGE_COMMITTED_PURCHASE_NO_UPFRONT	\$2.00	4
Compute	VM_SMALL	VM_SMALL_COMMITTED_HOUR	\$0.50	1
Compute	VM_MEDIUM	VM_MEDIUM_COMMITTED_HOUR	\$1.00	2
Compute	VM_LARGE	VM_LARGE_COMMITTED_HOUR	\$1.50	3
Compute	VM_XLARGE	VM_XLARGE_COMMITTED_HOUR	\$2.00	4
Compute	VM_SMALL	VM_SMALL_ON_DEMAND_HOUR	\$1.00	1
Compute	VM_MEDIUM	VM_MEDIUM_ON_DEMAND_HOUR	\$2.00	2
Compute	VM_LARGE	VM_LARGE_ON_DEMAND_HOUR	\$3.00	3
Compute	VM_XLARGE	VM_XLARGE_ON_DEMAND_HOUR	\$4.00	4

The above SKU Catalog shows that this provider only has 1 service that offers 4 virtual machine SKUs at various list rates, *commitment discount* rates, and normalization factors. Each SKU's normalization factor classifies its relative size to its *commitment discount* rate. Usage-based [commitment discounts](#)

with [commitment discount flexibility](#) can fully cover any combination of 1 or more SKUs where the sum of their normalization factor is less than or equal to the normalization factor of the *commitment discount*.

8.2.4.3. Scenario

- 1 no upfront *commitment discount* is purchased for 1 year (2023) for 1 VM_XLARGE which has a normalization factor of 8.
- 2 VM_MEDIUM resources run for 1 hour from 2023-01-01T00:00:00 to 2023-01-01T01:00:00 with a normalization factor of 4 for each.

8.2.4.4. Outcome

- 1 recurring, purchase record exists for 1 eligible "Normalized Hour" for a no upfront, *commitment discount* and incurs a \$2.00 [BilledCost](#).
- With *commitment discount flexibility*, 1 *commitment discount* for a VM_XLARGE covers 2 VM_MEDIUM resources within the corresponding [charge period](#) and incurs a \$2.00 total [EffectiveCost](#).
 - 1 *commitment discount* with a normalization factor of 8 covers 2 resources with normalization factors of 4 (i.e $4 + 4 = 8$).

[CSV Example](#)

8.3. Examples: Metadata

The following sections contain examples of metadata provided by a hypothetical FOCUS data provider called ACME to supply the required reference between the [FOCUS dataset](#) and the schema metadata. Provider implementations will vary on how the metadata is disseminated; however, the provider's chosen metadata delivery approach should be able to support the structure represented in this example.

In this example, the provider supports delivery of FOCUS data via file export to a data storage system. It uses JSON as the format for providing the metadata. The provider delivers data every 12 hours into a path structure described below:

Type of data	Path
Export location	/FOCUS
Metadata location	/FOCUS/metadata
Cost data location	/FOCUS/data

Here are some metadata examples for various scenarios:

8.3.1. Data Generator Metadata

8.3.1.1. Scenario

Acme provides metadata about the data generator as a part of their FOCUS data export. They provide the relevant data via the [Data Generator](#) schema object.

8.3.1.2. Supplied Metadata

Metadata can be provided at a location such as /FOCUS/metadata/data_generator.json.

The updated data generator related metadata could look like this:

```
{
  "DataGenerator": "Acme"
}
```

8.3.2. Deprecating Columns

8.3.2.1. Scenario

ACME has decided to deprecate columns prior to removal from their FOCUS data export. The column for deprecation is x_awesome_column3. The provider creates a new [Schema](#) object to represent the new schema, with a unique [SchemaId](#).

8.3.2.2. Supplied Metadata

Metadata can be provided at a location such as /FOCUS/metadata/schemas/schema-34567-abcde-34567-abcde-34567.json.

The updated schema related metadata could look like this:

```
{
  "SchemaId": "34567-abcde-34567-abcde-34567",
  "FocusVersion": "1.0",
  "CreationDate": "2024-03-02T12:01:03.083z",
  "ColumnDefinition": [
    {
      "ColumnName": "BillingAccountId",
      "DataType": "STRING",
      "StringMaxLength": 64,
      "StringEncoding": "UTF-8"
    },
    {
      "ColumnName": "BillingAccountName",
      "DataType": "STRING",
      "StringMaxLength": 64,
      "StringEncoding": "UTF-8"
    },
    {
      "ColumnName": "ChargePeriodStart",
      "DataType": "DATETIME"
    },
    {
      "ColumnName": "ChargePeriodEnd",
      "DataType": "DATETIME"
    },
    {
      "ColumnName": "BilledCost",
      "DataType": "DECIMAL",
      "NumericPrecision": 20,
      "NumberScale": 10
    }
  ]
}
```

```

        "NumberScale": 10
    },
    {
        "ColumnName": "EffectiveCost",
        "DataType": "DECIMAL",
        "NumericPrecision": 20,
        "NumberScale": 10
    },
    {
        "ColumnName": "Tags",
        "DataType": "JSON",
        "ProviderTagPrefixes": ["acme", "ac"]
    },
    {
        "ColumnName": "x_awesome_column1",
        "DataType": "STRING",
        "StringMaxLength": 64,
        "StringEncoding": "UTF-8"
    },
    {
        "ColumnName": "x_awesome_column2",
        "DataType": "DATETIME"
    },
    {
        "ColumnName": "x_awesome_column3",
        "DataType": "STRING",
        "StringMaxLength": 64,
        "StringEncoding": "UTF-8",
        "Deprecation": true
    }
]
}

```

For an example of how ACME ensures the schema metadata reference requirement is met see: [Schema Metadata to FOCUS Data Reference](#)

8.3.3. Renaming Columns

8.3.3.1. Scenario

ACME has decided to rename a column in their FOCUS data export. The column for rename is `x_awesome_column1` and will be renamed to `x_awesome_column_one`. The provider creates a new [Schema](#) object to represent the new schema, with a unique [SchemaId](#). After this schema definition is created if the data generator creates another schema, the `PreviousColumnName` is removed.

8.3.3.2. Supplied Metadata

Metadata can be provided at a location such as `/FOCUS/metadata/schemas/schema-34567-abcde-34567-abcde-34567.json`.

The updated schema related metadata for the schema where the rename took place could look like this:

```

{
  "SchemaId": "34567-abcde-34567-abcde-34567",
  "FocusVersion": "1.0",

```

```

"CreationDate": "2024-03-02T12:01:03.083z",
"ColumnDefinition": [
  {
    "ColumnName": "BillingAccountId",
    "DataType": "STRING",
    "StringMaxLength": 64,
    "StringEncoding": "UTF-8"
  },
  {
    "ColumnName": "BillingAccountName",
    "DataType": "STRING",
    "StringMaxLength": 64,
    "StringEncoding": "UTF-8"
  },
  {
    "ColumnName": "ChargePeriodStart",
    "DataType": "DATETIME"
  },
  {
    "ColumnName": "ChargePeriodEnd",
    "DataType": "DATETIME"
  },
  {
    "ColumnName": "BilledCost",
    "DataType": "DECIMAL",
    "NumericPrecision": 20,
    "NumberScale": 10
  },
  {
    "ColumnName": "EffectiveCost",
    "DataType": "DECIMAL",
    "NumericPrecision": 20,
    "NumberScale": 10
  },
  {
    "ColumnName": "Tags",
    "DataType": "JSON",
    "ProviderTagPrefixes": ["acme", "ac"]
  },
  {
    "ColumnName": "x_awesome_column_one",
    "DataType": "STRING",
    "StringMaxLength": 64,
    "StringEncoding": "UTF-8",
    "PreviousColumnName": "x_awesome_column1"
  },
  {
    "ColumnName": "x_awesome_column2",
    "DataType": "DATETIME"
  },
  {
    "ColumnName": "x_awesome_column3",
    "DataType": "STRING",
    "StringMaxLength": 64,
    "StringEncoding": "UTF-8",
    "Deprecation": true
  }
]
}

```


The subsequent new schema metadata after the rename could look like this:

```
{
  "SchemaId": "34567-abcde-34567-abcde-34567",
  "FocusVersion": "1.0",
  "CreationDate": "2024-03-02T12:01:03.083z",
  "ColumnDefinition": [
    {
      "ColumnName": "BillingAccountId",
      "DataType": "STRING",
      "StringMaxLength": 64,
      "StringEncoding": "UTF-8"
    },
    {
      "ColumnName": "BillingAccountName",
      "DataType": "STRING",
      "StringMaxLength": 64,
      "StringEncoding": "UTF-8"
    },
    {
      "ColumnName": "ChargePeriodStart",
      "DataType": "DATETIME"
    },
    {
      "ColumnName": "ChargePeriodEnd",
      "DataType": "DATETIME"
    },
    {
      "ColumnName": "BilledCost",
      "DataType": "DECIMAL",
      "NumericPrecision": 20,
      "NumberScale": 10
    },
    {
      "ColumnName": "EffectiveCost",
      "DataType": "DECIMAL",
      "NumericPrecision": 20,
      "NumberScale": 10
    },
    {
      "ColumnName": "Tags",
      "DataType": "JSON",
      "ProviderTagPrefixes": ["acme", "ac"]
    },
    {
      "ColumnName": "x_awesome_column_one",
      "DataType": "STRING",
      "StringMaxLength": 64,
      "StringEncoding": "UTF-8"
    },
    {
      "ColumnName": "x_awesome_column2",
      "DataType": "DATETIME"
    },
    {
      "ColumnName": "x_awesome_column3",
      "DataType": "STRING",
      "StringMaxLength": 64,
      "StringEncoding": "UTF-8",
      "Deprecation": true
    }
  ]
}
```

```
}  
]  
}
```

For an example of how ACME ensures the schema metadata reference requirement is met see:

[Schema Metadata to FOCUS Data Reference](#)

8.3.4. Schema Metadata

8.3.4.1. Scenario

ACME has only provided one [Schema](#) for their FOCUS data export. ACME provides a directory of schemas and each schema is a single file. Acme provides a file representing the schema for the data they provide.

8.3.4.2. Supplied Metadata

Metadata can be provided at a location such as /FOCUS/metadata/schemas/schema-1234-abcde-12345-abcde-12345.json.

The updated schema related metadata could look like this:

```

{
  "SchemaId": "1234-abcde-12345-abcde-12345",
  "FocusVersion": "1.0",
  "CreationDate": "2024-01-01T12:01:03.083z",
  "ColumnDefinition": [
    {
      "ColumnName": "BillingAccountId",
      "DataType": "STRING",
      "StringMaxLength": 64,
      "StringEncoding": "UTF-8"
    },
    {
      "ColumnName": "BillingAccountName",
      "DataType": "STRING",
      "StringMaxLength": 64,
      "StringEncoding": "UTF-8"
    },
    {
      "ColumnName": "ChargePeriodStart",
      "DataType": "DATETIME"
    },
    {
      "ColumnName": "ChargePeriodEnd",
      "DataType": "DATETIME"
    },
    {
      "ColumnName": "BilledCost",
      "DataType": "DECIMAL",
      "NumericPrecision": 20,
      "NumberScale": 10
    },
    {
      "ColumnName": "EffectiveCost",
      "DataType": "DECIMAL",
      "NumericPrecision": 20,
      "NumberScale": 10
    },
    {
      "ColumnName": "Tags",
      "DataType": "JSON",
      "ProviderTagPrefixes": ["acme", "ac"]
    }
  ]
}

```

8.3.5. Schema Metadata to FOCUS Data Reference

8.3.5.1. Scenario

ACME makes a change to the [Schema](#) of their data exports. For each FOCUS data export, ACME includes a metadata reference to the schema object. Because multiple files are provided in each export, Acme has elected to include a metadata file in each export folder that includes the FOCUS schema reference that applies to the data export files within that folder. When the schema changes, they include the new [Schema ID](#) in their export metadata file of the new folder.

8.3.5.2. Supplied Metadata

Metadata can be provided at a location such as `/FOCUS/data/export1-metadata.json`

The export metadata could look like this:

```
{
  "SchemaId": "1234-abcde-12345-abcde-12345",
  "data_location":
  [
    {
      "filepath": "/FOCUS/data/export1/export1-part1.csv",
      "total_bytes": 9010387,
      "total_rows": 4450
    },
    {
      "filepath": "/FOCUS/data/export1/export1-part2.csv",
      "total_bytes": 9010387,
      "total_rows": 4450
    },
    {
      "filepath": "/FOCUS/data/export1/export1-part3.csv",
      "total_bytes": 9010387,
      "total_rows": 4450
    },
    {
      "filepath": "/FOCUS/data/export1/export1-part4.csv",
      "total_bytes": 9010387,
      "total_rows": 4450
    }
  ]
}
```

New metadata can be provided at a location such as `/FOCUS/data/export2-metadata.json`.

The new export metadata could look like this:

```

{
  "SchemaId": "23456-abcde-23456-abcde-23456",
  "data_location":
  [
    {
      "filepath": "/FOCUS/data/export2/export2-part1.csv",
      "total_bytes": 9010387,
      "total_rows": 4450
    },
    {
      "filepath": "/FOCUS/data/export2/export2-part2.csv",
      "total_bytes": 9010387,
      "total_rows": 4450
    },
    {
      "filepath": "/FOCUS/data/export2/export2-part3.csv",
      "total_bytes": 9010387,
      "total_rows": 4450
    },
    {
      "filepath": "/FOCUS/data/export2/export2-part4.csv",
      "total_bytes": 9010387,
      "total_rows": 4450
    }
  ]
}

```

8.3.6. Data Changed by Data Generator Using Data Generator Version

8.3.6.1. Scenario

ACME specifies the optional metadata property [Data Generator Version](#) in their [Schema](#) object. They made a change to the [FOCUS dataset](#) they produce that does not adopt a new FOCUS Version, nor does it make a change the included columns, but does impact values in the data. This example illustrates that Data Generator Version changes are independent of column changes, however data generator version changes may include column changes.

The data generator creates a new schema object to represent the new schema. The data generator includes both the FOCUS Version and Data Generator Version in the schema object.

8.3.6.2. Supplied Metadata

Metadata can be provided at a location such as `/FOCUS/metadata/schemas/schema-56789-abcde-56789-abcde-56789.json`.

The updated schema related metadata could look like this:

```

{
  "SchemaId": "56789-abcde-56789-abcde-56789",
  "FocusVersion": "1.1",
  "DataGeneratorVersion": "2.4",
  "CreationDate": "2024-05-02T12:01:03.083z",
  "ColumnDefinition": [
    {
      "ColumnName": "BillingAccountId",
      "DataType": "STRING",
      "StringMaxLength": 64,
      "StringEncoding": "UTF-8"
    },
    {
      "ColumnName": "BillingAccountName",
      "DataType": "STRING",
      "StringMaxLength": 64,
      "StringEncoding": "UTF-8"
    },
    {
      "ColumnName": "ChargePeriodStart",
      "DataType": "DATETIME"
    },
    {
      "ColumnName": "ChargePeriodEnd",
      "DataType": "DATETIME"
    },
    {
      "ColumnName": "BilledCost",
      "DataType": "DECIMAL",
      "NumericPrecision": 20,
      "NumberScale": 10
    },
    {
      "ColumnName": "EffectiveCost",
      "DataType": "DECIMAL",
      "NumericPrecision": 20,
      "NumberScale": 10
    },
    {
      "ColumnName": "Tags",
      "DataType": "JSON",
      "DataGeneratorTagPrefixes": ["acme", "ac"]
    },
    {
      "ColumnName": "x_awesome_column1",
      "DataType": "STRING",
      "StringMaxLength": 64,
      "StringEncoding": "UTF-8"
    },
    {
      "ColumnName": "x_awesome_column2",
      "DataType": "DATETIME"
    }
  ]
}

```

For an example of how ACME ensures the schema metadata reference requirement is met see: [Schema Metadata to FOCUS Data Reference](#)

8.3.7. Adding New Columns

8.3.7.1. Scenario

ACME has decided to add additional columns to their FOCUS data export. The new columns are x_awesome_column1, x_awesome_column2, and x_awesome_column3. The provider creates a new [Schema](#) object to represent the new schema, this schema object has a unique [Schemaid](#). The subsequent data exports that use the new schema include the new schema's id as a reference to their corresponding schema object.

8.3.7.2. Supplied Metadata

Metadata can be provided at a location such as /FOCUS/metadata/schemas/schema-23456-abcde-23456-abcde-23456.json.

The updated schema related metadata could look like this:

```
{
  "Schemaid": "23456-abcde-23456-abcde-23456",
  "FocusVersion": "1.0",
  "CreationDate": "2024-02-02T12:01:03.083z",
  "ColumnDefinition": [
    {
      "ColumnName": "BillingAccountId",
      "DataType": "STRING",
      "StringMaxLength": 64,
      "StringEncoding": "UTF-8"
    },
    {
      "ColumnName": "BillingAccountName",
      "DataType": "STRING",
      "StringMaxLength": 64,
      "StringEncoding": "UTF-8"
    },
    {
      "ColumnName": "ChargePeriodStart",
      "DataType": "DATETIME"
    },
    {
      "ColumnName": "ChargePeriodEnd",
      "DataType": "DATETIME"
    },
    {
      "ColumnName": "BilledCost",
      "DataType": "DECIMAL",
      "NumericPrecision": 20,
      "NumberScale": 10
    },
    {
      "ColumnName": "EffectiveCost",
      "DataType": "DECIMAL",
      "NumericPrecision": 20,
      "NumberScale": 10
    },
    {
      "ColumnName": "Tags",
```

```

        "DataType": "JSON",
        "ProviderTagPrefixes": ["awecorp", "ac"]
    },
    {
        "ColumnName": "x_awesome_column1",
        "DataType": "STRING",
        "StringMaxLength": 64,
        "StringEncoding": "UTF-8"
    },
    {
        "ColumnName": "x_awesome_column2",
        "DataType": "DATETIME"
    },
    {
        "ColumnName": "x_awesome_column3",
        "DataType": "STRING",
        "StringMaxLength": 64,
        "StringEncoding": "UTF-8"
    }
]
}

```

For an example of how ACME ensures the schema metadata reference requirement is met see: [Schema Metadata to FOCUS Data Reference](#)

8.3.8. Removing Columns

8.3.8.1. Scenario

ACME has decided to remove columns from their FOCUS data export. The column removed is x_awesome_column3. The provider creates a new [Schema](#) object to represent the new schema, with a unique [Schemaid](#).

8.3.8.2. Supplied Metadata

Metadata can be provided at a location such as /FOCUS/metadata/schemas/schema-34567-abcde-34567-abcde-34567.json.

The updated schema related metadata could look like this:


```

{
  "SchemaId": "34567-abcde-34567-abcde-34567",
  "FocusVersion": "1.0",
  "CreationDate": "2024-03-02T12:01:03.083z",
  "ColumnDefinition": [
    {
      "ColumnName": "BillingAccountId",
      "DataType": "STRING",
      "StringMaxLength": 64,
      "StringEncoding": "UTF-8"
    },
    {
      "ColumnName": "BillingAccountName",
      "DataType": "STRING",
      "StringMaxLength": 64,
      "StringEncoding": "UTF-8"
    },
    {
      "ColumnName": "ChargePeriodStart",
      "DataType": "DATETIME"
    },
    {
      "ColumnName": "ChargePeriodEnd",
      "DataType": "DATETIME"
    },
    {
      "ColumnName": "BilledCost",
      "DataType": "DECIMAL",
      "NumericPrecision": 20,
      "NumberScale": 10
    },
    {
      "ColumnName": "EffectiveCost",
      "DataType": "DECIMAL",
      "NumericPrecision": 20,
      "NumberScale": 10
    },
    {
      "ColumnName": "Tags",
      "DataType": "JSON",
      "ProviderTagPrefixes": ["acme", "ac"]
    },
    {
      "ColumnName": "x_awesome_column1",
      "DataType": "STRING",
      "StringMaxLength": 64,
      "StringEncoding": "UTF-8"
    },
    {
      "ColumnName": "x_awesome_column2",
      "DataType": "DATETIME"
    }
  ]
}

```

For an example of how ACME ensures the schema metadata reference requirement is met see: [Schema Metadata to FOCUS Data Reference](#)

8.3.9. Changing Column Metadata

8.3.9.1. Scenario

ACME has decided to change the datatype of column x_awesome_column1 from a string to a number. ACME creates a new [Schema](#) object with the modification to x_awesome_column2.

8.3.9.2. Supplied Metadata

Metadata can be provided at a location such as /FOCUS/metadata/schemas/schema-67891-abcde-67891-abcde-67891.json.

The updated schema related metadata could look like this:

```

{
  "SchemaId": "67891-abcde-67891-abcde-67891",
  "FocusVersion": "1.0",
  "CreationDate": "2024-06-02T12:01:03.083z",
  "ColumnDefinition": [
    {
      "ColumnName": "BillingAccountId",
      "DataType": "STRING",
      "StringMaxLength": 64,
      "StringEncoding": "UTF-8"
    },
    {
      "ColumnName": "BillingAccountName",
      "DataType": "STRING",
      "StringMaxLength": 64,
      "StringEncoding": "UTF-8"
    },
    {
      "ColumnName": "ChargePeriodStart",
      "DataType": "DATETIME"
    },
    {
      "ColumnName": "ChargePeriodEnd",
      "DataType": "DATETIME"
    },
    {
      "ColumnName": "BilledCost",
      "DataType": "DECIMAL",
      "NumericPrecision": 20,
      "NumberScale": 10
    },
    {
      "ColumnName": "EffectiveCost",
      "DataType": "DECIMAL",
      "NumericPrecision": 20,
      "NumberScale": 10
    },
    {
      "ColumnName": "Tags",
      "DataType": "JSON",
      "ProviderTagPrefixes": ["acme", "ac"]
    },
    {
      "ColumnName": "x_awesome_column1",
      "DataType": "DECIMAL",
      "NumericPrecision": 20,
      "NumberScale": 10
    },
    {
      "ColumnName": "x_awesome_column2",
      "DataType": "DATETIME"
    }
  ]
}

```

For an example of how ACME ensures the schema metadata reference requirement is met see: [Schema Metadata to FOCUS Data Reference](#)

8.3.10. Provider Metadata Error Correction

8.3.10.1. Scenario

ACME has discovered that while their export includes the column `x_awesome_column3`, the [Schema](#) metadata does not include this column. In this case, the provider fixes the metadata in the existing schema object and does not need to create a new schema object. Reference metadata remains the same.

8.3.10.2. Supplied Metadata

Metadata can be provided at a location such as `/FOCUS/metadata/schemas/schema-34567-abcde-34567-abcde-34567.json`.

The updated schema related metadata could look like this:

```

{
  "SchemaId": "34567-abcde-34567-abcde-34567",
  "FocusVersion": "1.0",
  "CreationDate": "2024-03-02T12:01:03.083z",
  "ColumnDefinition": [
    {
      "ColumnName": "BillingAccountId",
      "DataType": "STRING",
      "StringMaxLength": 64,
      "StringEncoding": "UTF-8"
    },
    {
      "ColumnName": "BillingAccountName",
      "DataType": "STRING",
      "StringMaxLength": 64,
      "StringEncoding": "UTF-8"
    },
    {
      "ColumnName": "ChargePeriodStart",
      "DataType": "DATETIME"
    },
    {
      "ColumnName": "ChargePeriodEnd",
      "DataType": "DATETIME"
    },
    {
      "ColumnName": "BilledCost",
      "DataType": "DECIMAL",
      "NumericPrecision": 20,
      "NumberScale": 10
    },
    {
      "ColumnName": "EffectiveCost",
      "DataType": "DECIMAL",
      "NumericPrecision": 20,
      "NumberScale": 10
    },
    {
      "ColumnName": "Tags",
      "DataType": "JSON",
      "ProviderTagPrefixes": ["acme", "ac"]
    },
    {
      "ColumnName": "x_awesome_column1",
      "DataType": "STRING",
      "StringMaxLength": 64,
      "StringEncoding": "UTF-8"
    },
    {
      "ColumnName": "x_awesome_column2",
      "DataType": "STRING",
      "StringMaxLength": 64,
      "StringEncoding": "UTF-8"
    }
  ]
}

```

8.3.11. FOCUS Version Changed

8.3.11.1. Scenario

ACME's previous exports used FOCUS version 1.0. They are now going to adopt FOCUS version 1.1. It is required that they create a new schema metadata object which specifies the new FOCUS version via the [FOCUS Version](#) property—regardless of schema changes. In this example, the new FOCUS version adoption doesn't include columns changes. This is to illustrate that FOCUS version changes are independent of column changes, however, this scenario is unlikely.

8.3.11.2. Supplied Metadata

Metadata can be provided at a location such as `/FOCUS/metadata/schemas/schema-45678-abcde-45678-abcde-45678.json`.

The updated schema related metadata could look like this:

```

{
  "SchemaId": "45678-abcde-45678-abcde-45678",
  "FocusVersion": "1.1",
  "CreationDate": "2024-04-02T12:01:03.083z",
  "ColumnDefinition": [
    {
      "ColumnName": "BillingAccountId",
      "DataType": "STRING",
      "StringMaxLength": 64,
      "StringEncoding": "UTF-8"
    },
    {
      "ColumnName": "BillingAccountName",
      "DataType": "STRING",
      "StringMaxLength": 64,
      "StringEncoding": "UTF-8"
    },
    {
      "ColumnName": "ChargePeriodStart",
      "DataType": "DATETIME"
    },
    {
      "ColumnName": "ChargePeriodEnd",
      "DataType": "DATETIME"
    },
    {
      "ColumnName": "BilledCost",
      "DataType": "DECIMAL",
      "NumericPrecision": 20,
      "NumberScale": 10
    },
    {
      "ColumnName": "EffectiveCost",
      "DataType": "DECIMAL",
      "NumericPrecision": 20,
      "NumberScale": 10
    },
    {
      "ColumnName": "Tags",
      "DataType": "JSON",
      "ProviderTagPrefixes": ["acme", "ac"]
    },
    {
      "ColumnName": "x_awesome_column1",
      "DataType": "STRING",
      "StringMaxLength": 64,
      "StringEncoding": "UTF-8"
    },
    {
      "ColumnName": "x_awesome_column2",
      "DataType": "DATETIME"
    }
  ]
}

```

For an example of how ACME ensures the schema metadata reference requirement is met see: [Schema Metadata to FOCUS Data Reference](#)

8.3.12. FOCUS Version Changed by Data Generator Using Data

Generator Version

8.3.12.1. Scenario

ACME specifies the optional metadata property [Data Generator Version](#) in their [Schema](#) object. Their data generator version 2.2 supported FOCUS version 1.0. They are now going to adopt FOCUS Version 1.1 which requires that they update their Data Generator Version when updating the FOCUS Version. They create a new schema object designating that both properties have changed. In this example, the adoption of the new FOCUS version doesn't include additional columns. This is to illustrate that Data Generator Version can change independent of column changes; however, this scenario is unlikely.

The data generator creates a new schema object to represent the new schema. The data generator includes both the new FOCUS Version and Data Generator Version in the schema object.

8.3.12.2. Supplied Metadata

Metadata can be provided at a location such as `/FOCUS/metadata/schemas/schema-45678-abcde-45678-abcde-45678.json`.

The updated schema related metadata could look like this:


```

{
  "SchemaId": "45678-abcde-45678-abcde-45678",
  "FocusVersion": "1.1",
  "DataGeneratorVersion": "2.3",
  "name": "New Columns",
  "CreationDate": "2024-04-02T12:01:03.083z",
  "ColumnDefinition": [
    {
      "ColumnName": "BillingAccountId",
      "DataType": "STRING",
      "StringMaxLength": 64,
      "StringEncoding": "UTF-8"
    },
    {
      "ColumnName": "BillingAccountName",
      "DataType": "STRING",
      "StringMaxLength": 64,
      "StringEncoding": "UTF-8"
    },
    {
      "ColumnName": "ChargePeriodStart",
      "DataType": "DATETIME"
    },
    {
      "ColumnName": "ChargePeriodEnd",
      "DataType": "DATETIME"
    },
    {
      "ColumnName": "BilledCost",
      "DataType": "DECIMAL",
      "NumericPrecision": 20,
      "NumberScale": 10
    },
    {
      "ColumnName": "EffectiveCost",
      "DataType": "DECIMAL",
      "NumericPrecision": 20,
      "NumberScale": 10
    },
    {
      "ColumnName": "Tags",
      "DataType": "JSON",
      "ProviderTagPrefixes": ["acme", "ac"]
    },
    {
      "ColumnName": "x_awesome_column1",
      "DataType": "STRING",
      "StringMaxLength": 64,
      "StringEncoding": "UTF-8"
    },
    {
      "ColumnName": "x_awesome_column2",
      "DataType": "DATETIME"
    }
  ]
}

```

For reference, the prior schema object looked like this:

```

{
  "SchemaId": "34567-abcde-34567-abcde-34567",
  "FocusVersion": "1.0",
  "DataGeneratorVersion": "2.2",
  "CreationDate": "2024-04-02T12:01:03.083z",
  "ColumnDefinition": [
    {
      "ColumnName": "BillingAccountId",
      "DataType": "STRING",
      "StringMaxLength": 64,
      "StringEncoding": "UTF-8"
    },
    {
      "ColumnName": "BillingAccountName",
      "DataType": "STRING",
      "StringMaxLength": 64,
      "StringEncoding": "UTF-8"
    },
    {
      "ColumnName": "ChargePeriodStart",
      "DataType": "DATETIME"
    },
    {
      "ColumnName": "ChargePeriodEnd",
      "DataType": "DATETIME"
    },
    {
      "ColumnName": "BilledCost",
      "DataType": "DECIMAL",
      "NumericPrecision": 20,
      "NumberScale": 10
    },
    {
      "ColumnName": "EffectiveCost",
      "DataType": "DECIMAL",
      "NumericPrecision": 20,
      "NumberScale": 10
    },
    {
      "ColumnName": "Tags",
      "DataType": "JSON",
      "ProviderTagPrefixes": ["acme", "ac"]
    },
    {
      "ColumnName": "x_awesome_column1",
      "DataType": "STRING",
      "StringMaxLength": 64,
      "StringEncoding": "UTF-8"
    },
    {
      "ColumnName": "x_awesome_column2",
      "DataType": "DATETIME"
    }
  ]
}

```

For an example of how ACME ensures the schema metadata reference requirement is met see: [Schema Metadata to FOCUS Data Reference](#)

8.4. Examples: SaaS

The following section contains examples with how SaaS providers may implement the FOCUS specification. SaaS provider implementations will vary on the level of the detail available in their data, contract terms, purchasing options, and other factors.

8.4.1. Simple SaaS Agreements

Many SaaS providers provide simple contract terms, therefore don't need to support complex scenarios like spend commitments or pricing strategies in their billing data.

The scenarios described below illustrate how a FOCUS-compliant dataset should look for simple SaaS agreement scenarios (these scenarios may not be specific to SaaS agreements only).

8.4.1.1. Scenario A1: Invoice Up-front for a Purchase of a Service

ACME Corp allows its customers to purchase their service for a term (in this case, a year) for a \$10,000. ACME provides AwesomeCorp with a single invoice for their usage. ACME does not provide detailed cost and usage reports to AwesomeCorp throughout the Charge Period after the initial purchase.

Given that ACME does not charge based on or track usage, its usage details are irrelevant to this scenario.

[CSV Example](#)

Note the following details in the example dataset:

- The Charge Period is April 1st 2025 - April 1st 2026. The Billing Period is the month of April 2025 (when the licenses were ordered) and therefore will appear in the April invoice.
- A single charge representing the total payment for the 12-month agreement (\$10,000) is charged in the first invoice. BilledCost and EffectiveCost are realized in the same record since detailed usage records will not be provided during the 12-month period to realize amortized portions of this up-front payment.
- The single charge record does not include a List Unit Price, Pricing Quantity, or SKU-related information. Alternatively, the Pricing Quantity could have been set to 1, and the List Unit Price could be the same as the total charge.

8.4.1.2. Scenario A2: Invoice Up-front for a Quantity of a Service

ACME Corp offers its customer the ability to purchase a fixed quantity of licenses for their service. ACME provides AwesomeCorp with a single invoice for their usage. ACME does not provide detailed cost and usage reports to AwesomeCorp throughout the Charge Period after the initial purchase.

On April 1st, 2025, ACME executes a contract and invoices AwesomeCorp \$50,000 (Billed Cost) for a Charge Period of April 1st 2025 to April 1st 2026. As there is no negotiated discount, List Cost of the purchase is also \$50,000.

[CSV Example](#)

Note the following details in the example dataset:

- The Charge Period is April 1st 2025 to April 1st 2026. The Billing Period is the month of April 2025 (when the licenses were ordered) and therefore will appear in the April invoice.
- A single charge representing the total payment for the 12-month agreement is charged in the first invoice. Billed Cost and Effective Cost are both realized in the same record since detailed usage records will not be provided during the 12-month period to realize amortized portions of

this up-front payment.

- The single charge provided includes a ListUnitPrice for the licenses and a Pricing Quantity.

8.4.1.3. Scenario A3: Additional Purchase Records Provided in the SaaS Provider's FOCUS Data

On June 1st 2025 ACME provides the following records due to AwesomeCorp's \$1,000 mid-contract purchase of an additional 10 licenses for the same Charge Period (April 1st 2025 to April 1st 2026).

[CSV Example](#)

Note the following additional details in the example dataset:

- The Charge Period is still April 1st 2025 to April 1st 2026. The Billing Period is now the month of June 2025 (when the additional licenses were ordered) and therefore will appear in the June 2025 invoice.

8.4.1.4. Scenario B: Billed in Arrears for a Quantity of a Service

Similar to Scenario A above, ACME Corp offers its customer the ability to purchase their service with a fixed quantity of licenses. However, in Scenario B, ACME issues the invoice at the end of the usage period.

On April 1st, 2026, ACME invoices AwesomeCorp \$50,000 (Billed Cost) for the Charge Period of April 1st 2025 to April 1st 2026. As there is no negotiated discount, List Cost of the purchase is also \$50,000.

[CSV Example](#)

Note the following additional details in the example dataset:

- The Charge Period is April 1st 2025 to April 1st 2026. The Billing Period is now the month of March 2026 (since this charge is invoiced as of the last month of the Charge Period).

8.4.1.5. Scenario C: Simple SaaS Agreement with Monthly Billing

Like Scenario A2 above, ACME Corp offers its customers the ability to purchase their service with a fixed quantity of licenses. However, in Scenario C, ACME issues invoices at the end of each month (usage period). For this scenario, contract terms additionally include the following terms:

- ACME charges users monthly for the licenses that were consumed in that Billing Period
- The licenses are charged at \$20 per license per month

AwesomeCorp's consumption looks like this:

- In April 2025, AwesomeCorp uses 505 licenses
- In May 2025, AwesomeCorp uses 650 licenses
- In June 2025, AwesomeCorp uses 635 licenses

[CSV Example](#)

Note the following additional details in the example dataset:

- The Charge Period and Billing Period are April 1st, 2025, to May 1st, 2025, for the first month. Subsequent months increment the Charge Period and Billing Period by one month to match the month the charges are incurred.
- Billed Cost and Effective Cost are the same value since there is no up-front payment to amortize

8.4.2. SaaS Spend Agreements

Many SaaS providers support billing models that allow (or in some cases require) consumers to agree to an amount to spend over a period. In some cases, customers receive a negotiated discount for usage during that period in exchange for the spend agreement. Spend agreements can have different payment models like billing in arrears or pre-paid contracts and may impose minimum spend requirements for parts of the agreement.

The scenarios described below illustrate how a FOCUS-compliant dataset should look for various spend agreement scenarios.

8.4.2.1. Baseline Scenario

The following baseline conditions apply to the scenarios described below:

- AwesomeCorp has signed an agreement with SaaS provider Acme Co to use their database services
- On April 1 2025, AwesomeCorp agrees to spend \$1200 (post-discounts) in the upcoming 12-months
- AwesomeCorp receives a 20% negotiated discount in return for the commitment
- Acme Co calculates the spend counted against the agreements after discounts (like the negotiated discounts). Other providers may use the cost after discounts i.e., using List Cost for calculating the spend commitment.

8.4.2.2. Scenario A: Billed in arrears

For this scenario A, contract terms include the following terms in addition to the baseline scenario mentioned above:

- All charges will be billed in arrears at a monthly frequency

8.4.2.2.1. Scenario A1: Billed in arrears with no minimum spend requirement per month

For this scenario, contract terms additionally include the following terms:

- Committed spend can be used anytime within the 1-year term

AwesomeCorp's consumption looks like this:

- In the first month, AwesomeCorp uses \$48 of services (4 server hours). This usage has a List Cost of \$60 (before discounts)
- In the following 2 months, AwesomeCorp has some more usage
- For the final 9 months, AwesomeCorp does not use Acme services

[CSV Example](#)

Note the following details in the example dataset:

- A single charge representing the total unused amount from the 12-month agreement is charged during the final month of the 12-month term

8.4.2.2.2. Scenario A2: Billed in arrears with a minimum spend requirement per month

The spend agreement with Acme requires the customer to spend a minimum amount in each Billing Period (monthly). Unused fees are charged per Billing Period when the consumption is below this level (use-it or lose-it). For this scenario, contract terms additionally include the following terms:

- A minimum of \$60 needs to be spent each month

AwesomeCorp's consumption looks like this:

- In the first month, the AwesomeCorp uses \$48 of services (4 server hours). This usage has a List Cost of \$60 (before discounts). For this month, Acme charges \$12 (ListCost of \$15) for not meeting the monthly minimum
- In the following 2 months, AwesomeCorp has usage at or above the minimum requirement
- For the final 9 months, AwesomeCorp does not use Acme services

CSV Example

Note the following details in the example dataset:

- A monthly charge representing the unused minimum monthly amount is charged during months 4 through 11 of the 12-month term
- The final month has a charge that captures the overall unmet spend requirement for the 12-month contract. Alternatively, this could be provided as two charges, one for the unused portion of the final month, and one to capture the overall unmet spend requirement.

8.4.2.3. Scenario B: Prepaid contract

For this scenario B, contract terms include the following terms in addition to the baseline scenario mentioned above:

- The charges will be billed in arrears using monthly invoices

8.4.2.3.1. Scenario B1: Prepaid with no minimum spend requirement per month

Scenario B1 is similar to scenario A1 with the difference being that it's a pre-paid contract.

CSV Example

Note the following details in the example dataset:

- A purchase record for the initial \$1200 payment is present representing List, Billed, and Contracted cost of the purchase
- The charge for the unused amount has a \$0 BilledCost (since the total amount was billed with the prepayment). However, the charge captures the unused portion as an EffectiveCost.
- The unused charge rows apply to the entire Charge Period the contract was signed for.
- This scenario shows List Cost and Contracted Cost column double counting dynamic (described here in [ListCost](#) and [ContractedCost](#)) where either the [ChargeType](#) Purchase or Usage rows need to be excluded depending on the reporting scenario.

8.4.2.3.2. Scenario B2: Prepaid with a minimum spend requirement per month

Scenario B2 is similar to A2 with the difference being that it's a pre-paid contract.

CSV Example

Note the following details in the example dataset:

- A purchase record for the initial \$1200 payment is present representing List, Billed, and Contracted cost of the purchase

- The monthly charge for the unused amount has a \$0 BilledCost (since the total amount was billed with the prepayment). However, the charge captures the unused portion as an EffectiveCost.
- The final month has a charge that captures the overall unmet spend requirement for the 12-month contract. Alternatively, this could be provided as two charges, one for the unused portion of the final month, and one to capture the overall unmet spend requirement.
- This scenario shows List Cost and Contracted Cost column double counting dynamic (described here in [ListCost](#) and [ContractedCost](#)) where either the Purchase or Usage rows need to be excluded depending on the reporting scenario.

8.4.3. Virtual Currency Pricing Model

Many SaaS providers support pricing models that utilize virtual currencies such as credits, tokens, or points. Charges may be provided using a virtual currency, which can subsequently be converted to a national currency such as USD or EUR at an advertised or agreed-upon conversion rate.

The scenarios described below illustrate how a FOCUS-compliant dataset should look for various scenarios where a provider utilizes this pricing model.

8.4.3.1. Baseline Scenario

The following baseline conditions apply to the scenarios described below:

- AwesomeCorp has signed an agreement with SaaS provider Acme Co to use their services.
- Acme Co offers a virtual currency pricing model for their services and requires a purchase of virtual currency in advance of usage. Their denomination of virtual currency is called "tokens".
- Acme Co requires purchase of additional tokens in the event of usage exceeding purchased tokens.
- Acme Co publicly lists the cost of their tokens at \$2 per token.
- Acme Co treats token purchases as resources; therefore, charges for token purchases include values for ResourceId, ResourceName, and ResourceType.
- Acme Co publicly lists their usage to token rates. These rates are as follows:
 - 1 Q Widget Execution = 1 token
 - 1 Z Widget Execution = 2 tokens
 - 1 Workflow Operation = 3 tokens

8.4.3.2. Scenario A: Virtual Currency Not Offered at a Discount

For this scenario, contract terms include the following terms in addition to the baseline scenario mentioned above:

- Acme Corp offers no discount for purchased tokens.

8.4.3.3. Scenario A1: Purchase of Virtual Currency Without a Discount

For this scenario, the initial purchase of virtual currency is executed as follows:

- On April 1, 2025, AwesomeCorp agrees to purchase 100,000 tokens at \$2 per token for a total spend \$200,000. These tokens are only valid for 12 months.

[CSV Example](#)

Note the following details in the example dataset:

- The Charge Period is April 1st 2025 - April 1st 2026. The Billing Period is the month of April

- 2025 (when the tokens were purchased) and therefore will appear in the April invoice.
- Because Acme Co uses a virtual currency pricing model for usage and publishes their token price in terms of dollars and their usage cost in terms of tokens, their FOCUS dataset includes the columns PricingCurrency, PricingCurrencyContractedUnitPrice, PricingCurrencyEffectiveCost, and PricingCurrencyListUnitPrice.
- A single charge representing the total payment for the initial token purchase agreement (\$200,000) is charged in the first invoice.
 - ListCost, BilledCost, and ContractedCost of the purchase are all represented in this charge, however EffectiveCost is zero since the tokens are not yet consumed.
- PricingQuantity is set to the total tokens purchased.
- Because Awesome Corp is paying the list price, ListUnitPrice and ContractedUnitPrice are all set to the same value of \$2.

8.4.3.4. Scenario A2: Usage of Virtual Currency Purchased Without a Discount

Awesome Corp uses Acme's services consuming tokens as follows in the first day:

- 245 executions of Q Widget
- 5 executions of Z Widget
- 120 operations of Workflow

CSV Example

Note the following details in the example dataset:

- The Charge Period is April 1st 2025 - April 2nd 2025. The Billing Period is the month of April 2025.
- PricingCurrency for these usage charges reflects the per usage token price of the particular usage.
- PricingQuantity reflects the amount of usage of the PricingUnit for each charge and is equivalent to ConsumedQuantity. While relevant to this example, there are scenarios including tiered pricing where ConsumedQuantity and PricingQuantity may not be the same.
- Because Awesome Corp's usage includes no discount on usage to token rates, PricingCurrencyContractedUnitPrice and PricingCurrencyListUnitPrice are equivalent.

8.4.3.5. Scenario B: Virtual Currency Offered at a Discount

For this scenario, contract terms include the following terms in addition to the baseline scenario mentioned above:

- Acme Corp offers a discount for purchased tokens.

8.4.3.6. Scenario B1: Purchase of Virtual Currency at a Discount

For this scenario, the initial purchase of virtual currency is executed as follows:

- On April 1, 2025, AwesomeCorp agrees to purchase 100,000 tokens at discounted cost of \$1 per token for a total spend \$100,000. These tokens are only valid for 12 months.

CSV Example

Note the following details in the example dataset:

- The Charge Period is April 1st 2025 - April 1st 2026. The Billing Period is the month of April 2025 (when the tokens were purchased) and therefore will appear in the April invoice.
- Because Acme Co uses a virtual currency pricing model for usage and publishes their token

price in terms of dollars and their usage cost in terms of tokens, their FOCUS dataset includes the columns PricingCurrency, PricingCurrencyContractedUnitPrice, PricingCurrencyEffectiveCost, and PricingCurrencyListUnitPrice.

- A single charge representing the total payment for the initial token purchase agreement (\$100,000) is charged in the first invoice.
 - ListCost, BilledCost, and ContractedCost of the purchase are all represented in this charge, however EffectiveCost is zero, as required for prepaid purchases.
- PricingQuantity is set to the total tokens purchased.
- Because Awesome Corp is receiving a discount on the token price, the ListUnitPrice is set to \$2 and the ContractedUnitPrice is set to \$1. A ListCost of (\$200,000) and ContractedCost (\$100,000) reflect the cost of the tokens at the list price and contracted price respectively. The BilledCost is set to \$100,000 since this is the amount that Awesome Corp will be charged for the purchase of tokens.

8.4.3.7. Scenario B2: Usage of Virtual Currency Purchased at a Discount

Awesome Corp uses Acme's services, consuming tokens as follows in the first day:

- 245 executions of Q Widget
- 5 executions of Z Widget
- 120 operations of Workflow

CSV Example

Note the following details in the example dataset:

- PricingQuantity reflects the amount of usage of the PricingUnit for each charge and is equivalent to ConsumedQuantity. While relevant to this example, there are scenarios including tiered pricing where ConsumedQuantity and PricingQuantity may not be the same.
- Because Awesome Corp's usage includes no discount on usage to token rates, PricingCurrencyContractedUnitPrice and PricingCurrencyListUnitPrice are equivalent.

8.4.3.8. Scenario B3: Usage of Virtual Currency at a Modified Rate

Awesome Corp uses Acme's services consuming tokens as follows in the first day:

- 245 executions of Q Widget
- 5 executions of Z Widget
- 120 operations of Workflow

Additionally, Acme Co offers a modified usage to token ratio for one of their services as follows:

- 1 Workflow Operation = 2 tokens

CSV Example

Note the following details in the example dataset:

- Because of the modified rate for Workflow Operations, the PricingCurrencyContractedUnitPrice and PricingCurrencyListUnitPrice are different for this charge. The ContractedUnitPrice is set to \$1 and the ListUnitPrice is set to \$2.
- The PricingCurrencyEffectiveCost is 240 tokens for this charge, which is less than example B2 above due to the modified rate.
- ListCost reflects the cost of the charge at both the list cost of the tokens and the list rate for which the usage consumes tokens.

8.4.3.9. Scenario C: Handling Virtual Currency Usage Overages

For this scenario, Awesome Corp has exceeded their purchased tokens on October 1st 2025 by 1,500 tokens and Acme Co has charged them for the overage. The following conditions apply:

- Acme Co has charged Awesome Corp for the cost of tokens at the list price of \$2 per token, and this purchase is effective from April 1st 2025 to the date of the purchase, October 1st 2025.
- Awesome Corp purchases an additional 25,000 tokens to facilitate usage to the end of their contract. These tokens are valid from October 1st 2025 to April 1st 2026.

CSV Example

Note the following details in the example dataset:

- This example focuses on the purchases records only for the overage and additional purchases. Neither usage charges nor earlier purchases are not included in this example.
- The Charge Period for the Overage Purchase is April 1st 2025 - October 1st 2025. This is because the overage charge is to cover the period of time the overage token purchase is applicable to.
- The Charge Period for the Additional Purchase is October 1st 2025 - April 1st 2026. This is because the additional purchase is to cover the period of time to which the additional token purchase is applicable. Because end dates are exclusive, ChargePeriodEnd is April 1st 2026.

8.5. Grouping Constructs for Resources or Services

Providers natively support various constructs for grouping [resources](#) or [services](#). These grouping constructs are often used to mimic organizational structures, technical architectures, cost attribution/allocation and access management boundaries, or other customer-specific structures based on requirements.

Providers may support multiple levels of resource or service grouping mechanisms. FOCUS supports two distinct levels of groupings that are commonly needed for FinOps capabilities like chargeback, invoice reconciliation and cost allocation.

- [Billing account](#): A mandatory container for *resources* or *services* that are billed together in an invoice. *Billing accounts* are commonly used for scenarios like grouping based on organizational constructs, invoice reconciliation and cost allocation strategies.
- [Sub account](#): An optional provider-supported construct for organizing *resources* and *services* connected to a *billing account*. *Sub accounts* are commonly used for scenarios like grouping based on organizational constructs, access management needs and cost allocation strategies. *Sub accounts* must be associated with a *billing account* as they do not receive invoices.

The table below highlights key properties of the two grouping constructs supported by FOCUS.

Property	Billing account	Sub account
Requirement level	Mandatory	Optional
Receives an invoice?	Yes	No
Invoiced at	Self	Associated billing account
Examples	AWS: Management Account* GCP: Billing Account Azure MCA: Billing Profile Snowflake: Organizational Account	AWS: Member Account GCP: Project Azure MCA: Subscription Snowflake: Account

* For organizations that have multiple AWS Member Accounts within an AWS Organization, consolidated billing is enabled by default and invoices are received at Management Account level. A Member Account can be removed from AWS consolidated billing whereby the removed account receives independent invoices and is responsible for payments.

8.6. Origination of Cost Data

Cost data presented in the billing datasets originates from various sources depending on the purchasing mechanism. There are at least 3 different pieces of information that are important for understanding where cost originated from.

- **Provider:** The entity that made the *resources* or *services* available for purchase.
- **Publisher:** The entity that produced the *resources* or *services* that were purchased.
- **Invoice Issuer:** The entity responsible for invoicing for the *resources* or *services* consumed.

The value for each of these may be different depending on the various purchasing scenarios for *resources* or *services*. Use cases for purchasing direct, via a Managed Service Provider (MSP), via a cloud marketplace, and from internal service offerings were considered. The table below presents a few scenarios to show how the value for each dimension may change based on the purchasing scenario.

#	Scenario	Provider	Publisher	Invoice Issuer
1.1	Purchasing cloud services directly from cloud provider	Cloud service provider	Cloud service provider	Cloud service provider
1.2	Purchasing cloud services from the cloud provider where the cloud region is operated by a 3rd party	Cloud service provider	Cloud service provider	Entity operating the region for the cloud service provider
2.1	Purchasing cloud services via MSP	Managed Service Provider	Cloud service provider	Managed Service Provider
2.2	Purchasing cloud-agnostic resources or services built/sold by an MSP	Managed Service Provider	Managed Service Provider	Managed Service Provider
2.3	Purchasing labor services from managed service provider	Managed Service Provider	Managed Service Provider	Managed Service Provider
3.1	Purchasing a cloud marketplace offering that runs on the cloud provider	Cloud service provider	Company building the software or services (Cloud service provider OR third-party software or services company)	Cloud service provider
3.2	Purchasing a cloud marketplace offering that is not running directly on your cloud infrastructure (e.g., SaaS product, Professional Services)	Cloud service provider	Company producing the SaaS or services product	Cloud service provider
3.3	Purchasing a SaaS product that is not directly running on your cloud infrastructure from a 3rd party reseller managed cloud marketplace	Cloud service provider	SaaS provider	Reseller
4.1	Purchasing SaaS software directly from provider	SaaS provider	SaaS provider	SaaS provider
4.2	Purchasing SaaS software that additionally runs on your cloud resources (in addition to #4.1)	Cloud service provider	Cloud service provider	Cloud service provider

#	Scenario	Provider	Publisher	Invoice Issuer
5.1	Purchasing internal infrastructure or services offerings running on-premise	Internal product name	Internal product name	Internal product name
5.2	Purchasing internal infrastructure or services offerings running on cloud	Internal product name	Internal product name	Internal product name
5.3	Associated software license cost for use on an on-premise infrastructure platform (Where license cost is presented separately in cost data)	Internal product name	Company producing the software	Internal product name