

Administrator Guide

MAINTENANCE RELEASE 85

Copyright notice & disclaimers

Copyright © 2000–2020 PortaOne, Inc. All rights reserved

PortaBilling® Administrator Guide, June 2020

Maintenance Release 85

V1.85.06

Please address your comments and suggestions to: Sales Department,
PortaOne, Inc. Suite #408, 2963 Glen Drive, Coquitlam BC V3B 2P7
Canada.

Changes may be made periodically to the information in this publication. The changes will be incorporated in new editions of the guide. The software described in this document is furnished under a license agreement, and may be used or copied only in accordance with the terms thereof. It is against the law to copy the software on any other medium, except as specifically provided for in the license agreement. The licensee may make one copy of the software for backup purposes. No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopied, recorded or otherwise, without the prior written permission of PortaOne Inc.

The software license and limited warranty for the accompanying products are set forth in the information packet supplied with the product, and are incorporated herein by this reference. If you cannot locate the software license, contact your PortaOne representative for a copy.

All product names mentioned in this manual are for identification purposes only, and are either trademarks or registered trademarks of their respective owners.

Table of Contents

Preface	7
Hardware and software requirements	8
Installation	9
What is new in maintenance release 85?	10
1. System architecture	11
Overview	12
Components	14
PortaBilling® web interface	17
PortaBilling® performance	18
High availability	23
2. System concepts	28
Overview	29
Service types and services	30
Service provisioning	36
Products	37
Add-on products	39
Retail customers	46
Account and account types	54
Customer and account statuses	61
Customer business models and account roles	70
Distributors	74
Resellers	79
Representatives	89
Vendors	94
Nodes	96
Virtual environments	97
Billing configuration	99
Billing entities hierarchy	101
3. Real-time charging	102
Destinations, rates and tariffs	103
Override tariff	109
Volume discount plans	111
Transfer of funds and airtime among end users	137
Billing for time-based services	140
Flexible definition of the party to charge	148
Billing for quantity-based services	149
Overdraft protection	150
4. Recurring charging	157
Subscriptions	158
Bundle promotions	177
Billing based on measured resource consumption	183
5. Voice services	189
Rate upload	190

xDR calculation	191
xDR recalculation	196
Call routing	200
Number translation	220
Favorite number billing.....	226
US inter / intra state billing	226
DID provisioning and billing	227
Registration for PINless dialing service via SMS.....	243
Origin-based pricing for European destinations.....	244
OLI-based billing	247
6. MNO / MVNO billing	250
Integration with mobile networks.....	251
Mobile Virtual Network Operator.....	254
SMS routing with real-time HLR lookup	255
Real-time billing for roaming services.....	257
SIM card inventory.....	261
Mobile account batch generation for SIM card distribution.....	262
Local number portability and other special cases	265
Integration with NetNumber	269
Roaming identification in xDRs	270
Twin card / data card services.....	271
7. LTE service provisioning	272
LTE service provisioning and billing.....	273
VoLTE.....	279
Verified devices.....	281
8. VoIP services.....	284
Service policies.....	285
Limiting per-account simultaneous calls on a product.....	287
Customer Premises Equipment (CPE) provisioning and inventory.....	289
Bandwidth-aware limitation on concurrent calls permitted.....	291
Active call monitoring.....	292
Using different protocols (SIP and H323)	292
9. Fraud protection tools	294
Tools for prevention of VoIP fraud.....	295
Protection from DoS attacks	301
Toll fraud prevention.....	302
Fraud traffic profiles	303
Binding a user session to an IP address	309
10. Internet provisioning and charging.....	310
Broadband Internet access service	311
Internet service provisioning for ADSL users via Cisco 7000 / 7200 series routers.....	315
Fiber to the home service provisioning via Calix ONT.....	318
WiMAX services.....	320
Captive portal for WiFi hotspots.....	323

11. IPTV provisioning	327
IPTV.....	328
12. xDR mediation and import	330
xDR Mediation concept.....	331
13. Billing / invoicing	337
Multi-currency support.....	338
Invoicing concept.....	338
Collection policy.....	345
Initial invoices.....	352
Midterm invoices.....	353
Out-of-turn invoices.....	353
Balance adjustment.....	354
Processing taxes.....	359
Payments.....	377
14. Web interface administration	390
Multi-language support.....	391
Access levels.....	393
Role-based access control.....	400
Limits management for administrative users.....	403
Quick forms.....	407
Custom fields.....	407
Online web signup.....	407
Personal data protection. GDPR compliance.....	408
White labeling.....	415
15. Operations and integrations	421
PortaOne iPaaS.....	422
Integration with the Microsoft Active Directory.....	424
Integration with Salesforce CRM.....	427
Invoice data export to SAP Business One ERP.....	430
Troubleshooting Tools.....	431
Password protection.....	433
16. Statistics and reporting	436
Date and time information.....	437
xDR summary.....	438
Custom reports.....	440
PortaBilling server health monitoring.....	445
VoIP network performance statistics.....	447
17. Appendices	452
APPENDIX A. Supported RADIUS RFCs.....	453
APPENDIX B. Supported Diameter RFCs.....	453
APPENDIX C. Change-of-Authorization (CoA) and Disconnect Message (DM) packets supported by PortaBilling®.....	453
APPENDIX D. The Number_Portability table description.....	454
APPENDIX E. Billing diagrams.....	455

APPENDIX F. Supported online payment processors.....458
APPENDIX G. Personal data that can be stored in PortaBilling®.....459
APPENDIX H. Glossary / list of abbreviations462

Preface

This document provides PortaBilling® users with the most common examples and guidelines for setting up a VoIP network or enhancing a network with services such as IPTV.

Where to get the latest version of this guide

The hard copy of this guide is updated upon major releases only, and does not always contain the latest material on enhancements that occur in-between minor releases. The online copy of this guide is always up to date, and integrates the latest changes to the product. You can access the latest copy of this guide at www.portaone.com/support/documentation/.

Conventions

This publication uses the following conventions:

- Commands and keywords are given in **boldface**.
- Terminal sessions, console screens, or system file names are displayed in `fixed width font`.



Exclamation mark draws your attention to important actions that must be taken for proper configuration.

NOTE: Notes contain additional information to supplement or accentuate important points in the text.



Timesaver means that you can save time by taking the action described here.



Tips provide information that might help you solve a problem.



Archivist explains how the feature worked in previous releases.



Gear points out that this feature must be enabled on the Configuration server.

Trademarks and copyrights

PortaBilling®, PortaSIP® and PortaSwitch® are registered trademarks of PortaOne, Inc.

Hardware and software requirements

Server system recommendations

- Four (4) Linux Servers for the PortaBilling®:
 - Real-time billing engine (RADIUS).
 - Main database server.
 - Server for web-interfaces (web interface for administrators and self-care portal, XML API services, periodic tasks, invoicing) and replica database.
 - Configuration management (this server is also used for monitoring the other servers and collecting and processing log information).
- A minimum of 250 GB of available disk space on each billing server. On the average, 200,000 CDRs take up about 1 GB of disk space (this includes database files, indexes and binary logs, raw RADIUS detail files, billing engine log files and other related information), plus you need to reserve an amount of free space roughly equal to the projected database size for performing operations such as backup. RAID is recommended in order to improve performance and reliability.

The Configuration server requires about 500 GB of free disk space if you are using PortaSwitch® (where log files from the PortaSIP® dispatching and processing nodes need to be processed), and 250 GB of free disk space if you are using the PortaBilling® standalone.

- A 64-bit processor (Xeon, Opteron).
- At least 8 GB of RAM, 16 GB recommended.
- Two Gigabit network ports (can be two separate network adapters or a single dual-port adapter).

For additional details and configuration advice, see the *Hardware Recommendations* topic on our website:

<http://www.portaone.com/support/hw-requirements/>

For information about whether particular hardware is supported by Oracle Enterprise Linux used as the operating system in PortaSwitch®, consult the related document on the Oracle or RedHat website:

<https://hardware.redhat.com/>.

Client system recommendations

- **OS:** MS Windows 7 or above, Linux/BSD, macOS Sierra (version 10.12) or above.
- **Web browser:**
 - Google Chrome 76 or above, Chromium 76 or above, Mozilla Firefox 68 or above.
 - JavaScript and cookies must be enabled.
- **Spreadsheet processor:** MS Excel, OpenOffice Calc, LibreOffice Calc, Google Sheets.
- **Display settings:** a minimum screen resolution of 1366 × 768.

Installation

PortaSwitch® installation ISO files contain everything required for installing Oracle Enterprise Linux (64-bit version), PortaSwitch® and the supplementary packages that are necessary for convenient system administration and maintenance.

After the installation is complete you will assign roles (e.g. RADIUS, web interface, etc.) to individual servers using the Configuration server tool – this will automatically enable the required components of PortaSwitch® software on each server.

PortaSwitch® installation and configuration are automated and integrated within the main installation process. This allows you to install a completely functional PortaSwitch® environment from scratch in less than one hour!

For detailed installation instructions, please refer to the [PortaBilling® Installation Guide](#).

What is new in maintenance release 85?

Added:

- The **Mobile account batch generation for SIM card distribution** chapter.
- The **Integration with the Microsoft Active Directory** chapter.

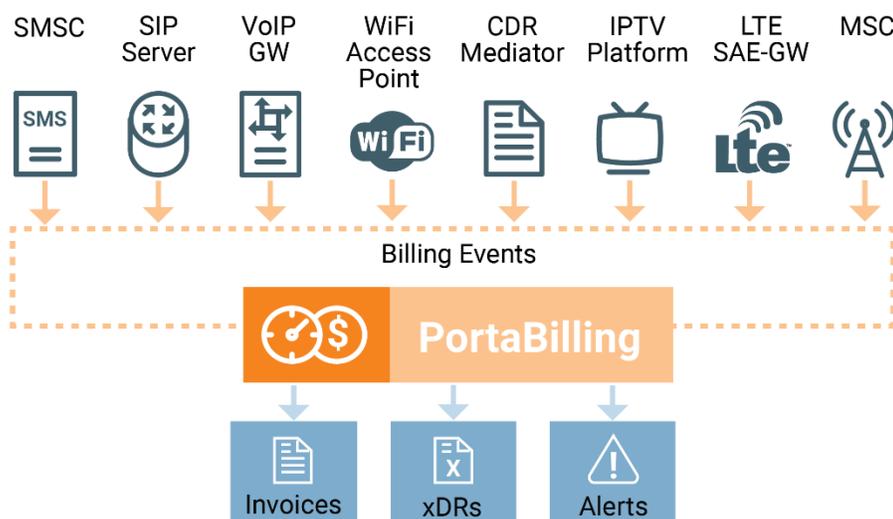
Updated:

- The **Subscription fee** chapter.
- The **Multi-language support** chapter.
- The **Identical pre-condition** chapter.

1 ■ System architecture

Overview

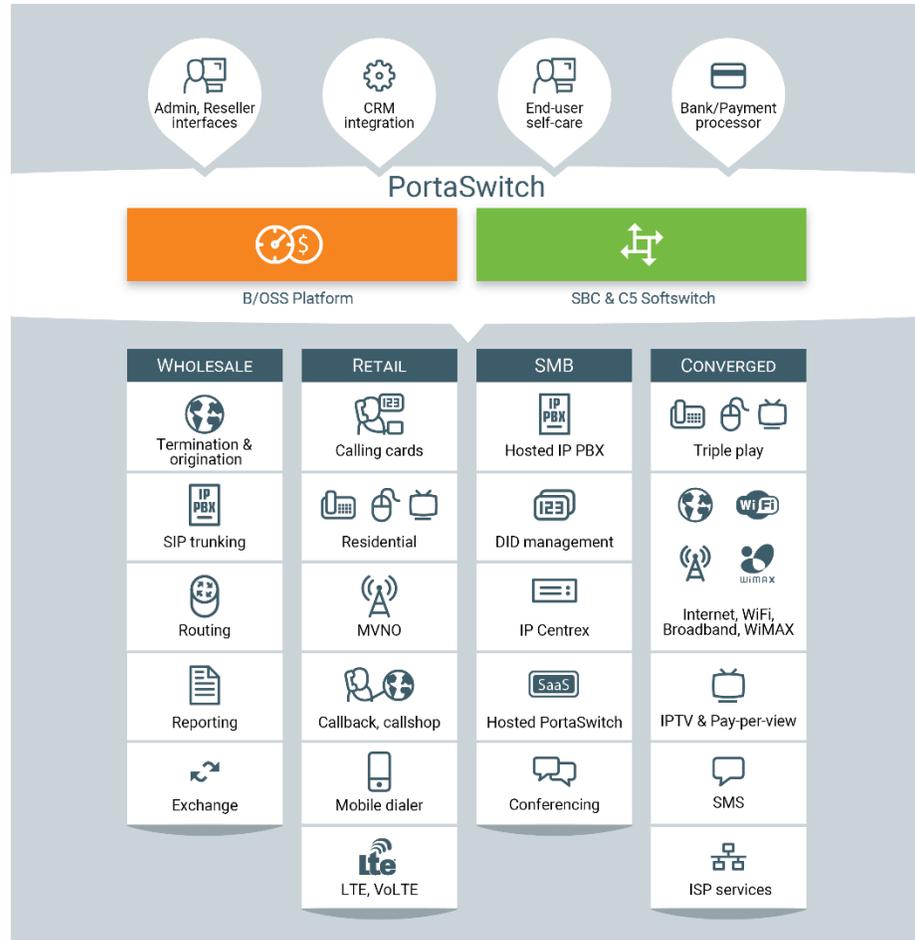
PortaBilling® is a carrier-grade converged billing / provisioning system for communication service providers. It communicates with elements of your network (such as VoIP gateways or WiFi/DSL access points), provides these nodes with authentication or authorization (determining whether a customer should be admitted and provided with a service), and gathers billing events, i.e. data about services rendered to your customers. Based on this information, it performs rating for the services, creates transaction records (also called xDRs - eXtensible Detail Records¹), and modifies customers' balances accordingly.



All this happens in real time, so the billing data is updated as soon as a session is completed (e.g. the customer hangs up his phone, or an SMS message has been sent). PortaBilling® provides a unified platform for multiple services, which allows you to use it to charge clients for their voice calls, messages, and data transfer, thus effectively deploying triple-play on your network.

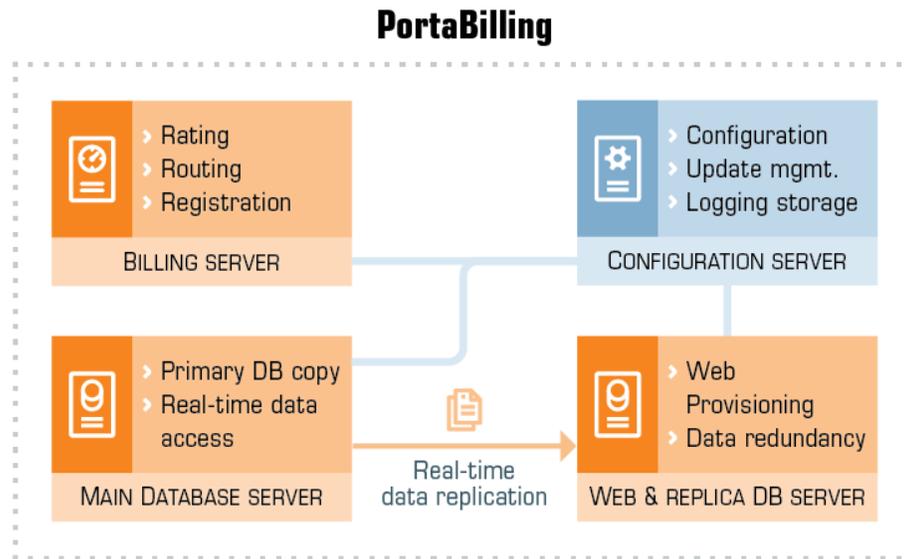
PortaBilling® will act as the nerve center of your network. After you have entered information about your services, rates, customers and so on via the web interface, PortaBilling® will communicate in real time with elements of your network to supply information regarding which customers the service should be provided to (and which not), as well as exactly how it should be provided. Customers whose balance has run out will be disconnected immediately after exceeding the maximum session duration and (since billing happens in real time) funds will be withdrawn from their account and service denied if they make another attempt to use the service.

¹ For xDRs for telephony services, the previous term CDRs (Call Details Records) is often used.

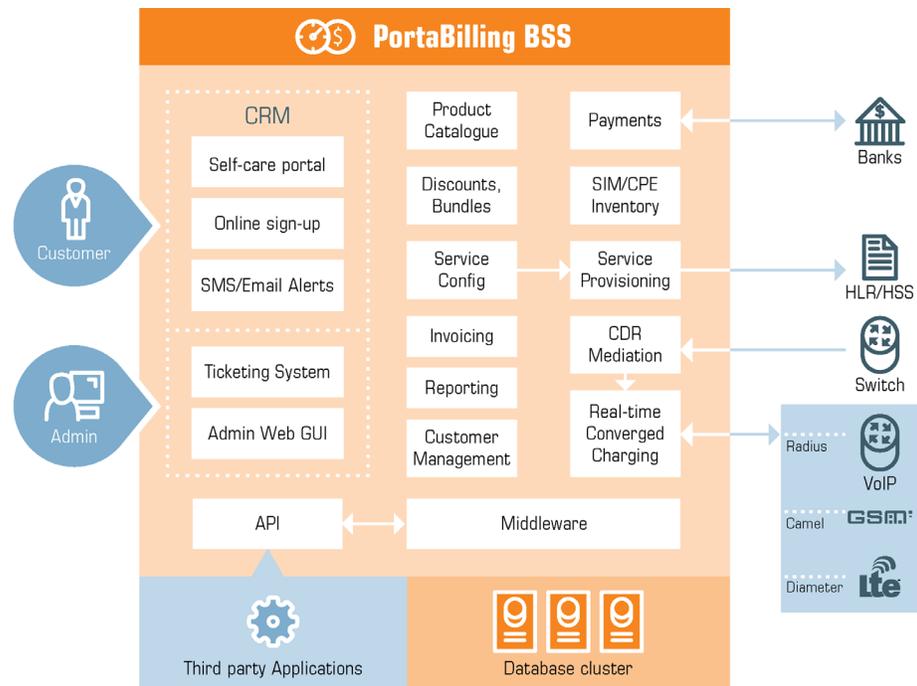


As shown in the diagram above, multiple types of network equipment can work in conjunction with PortaBilling®. The easiest and best method of integration is via the RADIUS protocol. However, if a certain type of switch or server does not support RADIUS, integration can be done by other means, for instance XML API.

Components



PortaBilling® consists of four main logical components: the main database server (lower left), the billing (online charging) server (upper left), the admin/web and replica database server (bottom right), and the configuration/log collector server (upper right). Typically, the billing, web and two database servers are installed on three separate physical computers so as to provide data redundancy and load sharing. Although a “shared” combination (where some servers, e.g. billing engine and main database run on a single physical computer) is also possible and supported, this is recommended only for evaluation or demo systems.



The Configuration server must always reside on a separate physical server. It alters the software configuration and version of the other servers, and must remain unaffected by the changes thus made. This allows a rollback to the previous stable version of the configuration or software release.

Online charging server (OCS)

The OCS server hosts the billing engine which communicates with the nodes via the RADIUS protocol. It can also process Diameter requests that are received from nodes provisioned for services via the Diameter protocol.

The billing engine provides:

- **Authentication** – It tells the node whether the subscriber (identified by phone number, PIN, IP, or the like) is allowed to use a specific service (e.g. voice calls or wireless Internet access), and returns attributes such as current balance.
- **Authorization** – It tells the node whether the subscriber (identified by phone number, PIN, IP, or the like) is allowed to initiate a session with specific parameters within a service (e.g. calling a particular phone number), and returns session attributes such as the maximum allowed session duration or the allowed amount of bandwidth.
- **Interim (also called keep-alive) accounting processing** – This is used for services like Internet access, when nodes constantly update the billing about usage for the current session.

- **Accounting processing** – Based on information from the gateways, it bills the session and writes transaction records to the database.

Main database server

The main database server hosts the primary copy of the database. This database is used for all real-time data access (e.g. during account authentication) and for all data modifications.

Web and replica database server

Web server

The web server hosts the following:

- Web interface, consisting of:
 - Admin interface.
 - Account self-care.
 - Customer self-care.
 - Reseller's helpdesk (self-care; customer care).
 - Distributor self-care.
 - Vendor self-care.
 - Representative self-care.
 - Online web signup interface.
- Scripts for generating invoices, calculating various statistics and producing reports.
- Optional IVR (TFTP) server.
- XML API (SOAP) interface for integration with external applications.

Replica database server

The replica database server hosts the second copy of the database. This database is used for all data mining activities (e.g. calculating a summary for the invoice or producing reports) and can be used to restore billing data in the event that the main database is lost.

Configuration server

This server stores the information about all the other servers in the installation plus service configuration details about each of them. See more details about the configuration management framework in the *Centralized Configuration Management* chapter of the **PortaSwitch Architecture and Concepts** guide.

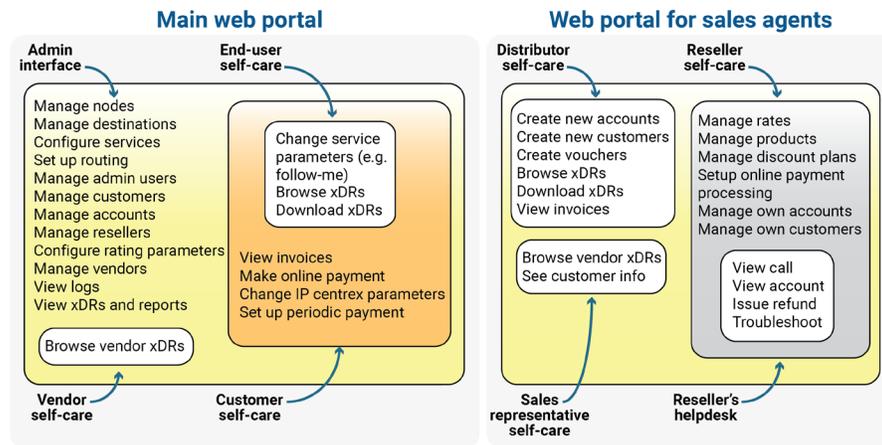
The Configuration server monitors the status of all the other servers. Data such as available disk space, load, or service status can be viewed via the web interface. If a problem is detected, the PortaOne support team will be notified immediately.

The Configuration server also collects log data, such as SIP communication logs from PortaSIP®, or billing engine logs from billing nodes. These are then stored, processed and indexed, so that all entries relating to a particular session can be retrieved instantly. This reduces the amount of disk space needed on the production servers, as well as the load on them.

Log data for the most recent time period is kept in the original format. Older data is archived to save disk space, while log files which are outdated are erased.

PortaBilling® web interface

Different operations are available for different types of users who access the system (also, the abilities of an individual user can be adjusted via his access level settings). In general, administrators (PortaBilling® users) can potentially perform all possible configuration tasks, while others are limited to their own entity and related objects only. For instance, a reseller can access any of his subcustomers or their accounts, while a subcustomer can only access information about himself and his accounts on his self-care pages. The self-care interface imposes multiple security-related limitations on a user’s activity, e.g. a subcustomer can view his own balance on the self-care pages, but cannot modify it.



For improved security, the PortaBilling® web interface is divided into several separate areas, each serving its own purpose. Thus a login name and password which allow a user to access the customer self-care interface cannot be used to log in to the admin interface; the login name and

password for the account's (end-user's) self-care interface are not applicable to customer self-care interface; and so forth.

PortaBilling® performance

When buying a new car, you are obviously interested in performance parameters such as “how fast can it go?” The maximum speed you can achieve depends, of course, on your car's technical parameters, such as engine horsepower strength or total car body weight. In general, the more powerful the car you buy, the faster it will go. However, the actual speed of the car at any given moment depends on many other criteria too, such as how many passengers are in the car, how much weight is being carried in its trunk, what are the road conditions (paved highway vs. gravel road), etc.

The situation with billing systems is quite similar. Typical questions are: ‘How many customers can I have?’, ‘How many concurrent calls can the system handle?’ or ‘How many messages or call minutes per month can I transport on my network?’. The actual numbers depend not only on PortaBilling® functionality, your hardware's specification, but also on your billing configuration, traffic patterns and types of services provided.

PortaBilling® communicates with the network nodes via the RADIUS and Diameter protocols for session authorization and rating. Let's take a more detailed look at its performance for either protocol.

Note that the performance results are given for a single billing server with 32 active processing engines running on the **recommended hardware**. Performance results will be different for different number of processing engines. The performance of PortaBilling® cluster is higher than that of a single billing server.

PortaBilling® RADIUS performance

Each call attempt or completed call on your network generates several RADIUS requests, with the total number of requests and their type dependent upon the kind of service provided. For services such as IP Centrex, which involve call authorization, route calculation and xDR processing, PortaBilling® can handle about 90 CAPS (call attempts per second). This means that each second 90 users can begin a new phone call on your network. For postpaid wholesale services in which PortaBilling® only processes the xDRs, the number of call attempts per second increases to about 230. The exact figures will depend upon the complexity of your billing configuration (e.g. how many rate plans, individual rates in a rate plan or carriers you have). Since it is not possible to know which

combination of different services you will provide on your network beforehand, it only makes sense to refer to a conservative **estimate** of the PortaBilling® processing speed, which has been calculated at **90 session initiations per second**.

How many simultaneous (concurrent) calls does that translate to?

This primarily depends on the average call duration and call success rate. Assuming an call processing speed of 90 CAPS, an average call duration of 5 minutes and a call success rate of 50% (industry norms), 50% of the 90 CAPS will succeed. This means that 45 calls will become connected, and the same amount of previously connected calls will become disconnected. Since average call duration is 300 seconds, approximately $45 * 300 = 13,500$ calls will be in a “connected” state at all times. If your ASR changes, it will immediately impact the number of concurrent calls, e.g.:

- ASR 25%: $(90 * 25\%) * 300 = 22.5 * 300 = 6,750$ concurrent calls.
- ASR 60%: $(90 * 60\%) * 300 = 54 * 300 = 16,200$ concurrent calls.

How many minutes of monthly traffic can PortaBilling® support?

Again, this depends mainly on your traffic patterns (the ratio between the duration of your peak and off-peak hours and the numbers of calls during peak and / or off-peak hours) and your call success rate. We can make estimates based on typical traffic patterns (as indicated below) that occur during daily peak hours, medium load hours and off-peak hours. We should also bear in mind that your network must be able to handle all kinds of traffic – off-peak hour, peak hour and extraordinary traffic splashes (e.g. New Year’s Eve) where there are 2 to 3 times more call attempts than during normal peak hours. Thus, the total performance capacity must be greater than the maximum anticipated BHCA (busy hours call attempts).

Using the same “industry norm” figures as above (50% ASR, 5 mins ALOC), PortaBilling® can process $(90 * 50\%) * 5 * 3600 = 810,000$ minutes during peak hours. We can estimate that your total amount of daily traffic will be about 6 times the amount of traffic during a single peak hour, i.e. $810,000 * 6 = 4,860,000$ minutes per day. Then, multiplying the daily amount by the number of “business” days in a month, we get $4,860,000 * 20 = 97,200,000$ minutes.

This is why, in our estimate (to be on the safe side and allow for some reserve capacity), PortaBilling® is poised to handle monthly traffic of over **95 million minutes**.



The results are calculated based on the **maximum** server’s capacity, which is 90 CAPS. However, we strongly advise you to reserve some capacity of the billing server when you plan your infrastructure. This is required to handle traffic splashes and the traffic increase when your business grows.

PortaBilling Diameter performance

In mobile networks, the core components communicate with PortaBilling® via the Diameter (Gy) interface for Internet sessions and via Diameter (Ro) interface for voice calls and SMS. Therefore, let’s consider PortaBilling®’s performance for each, separately.



Note that the performance data given below is for informational purposes only. It is measured for every service separately. Your real performance highly depends on your server specs and the number of servers deployed.

Voice calls

Similar to call processing via RADIUS, every call attempt in a mobile network results in several Diameter requests. Assuming that a call attempt includes call authorization, session update and xDR generation, PortaBilling® can handle about 140 CAPS (call attempts per second). This means that 140 users can begin a new phone call on your network every second. The exact figures highly depend on your billing configuration: how many rate plans you have, how many rates there are in your rate plans, what volume discounts you apply, etc.

To calculate how many concurrent calls this translates into, we must consider such call parameters as average success rate (ASR) and average length of call (ALOC). The ASR varies from country to country. For our estimation we’ll take the average ASR value as 60% and an average call length as 3 minutes. Given the processing speed of 140 CAPS, this means that $140 * 60\% = 84$ calls become connected and the same number of previously connected calls become disconnected. With the average call length of 3 minutes, this results into $84 * 180 = 15.200$ concurrent calls.

The table below shows how the number of concurrent calls differs depending on the ASR and ALOC values:

CAPS	ASR, %	ALOC, sec	Number of concurrent calls
140	60	180	15,200
140	60	90	7,560
140	90	180	22,600

The amount of monthly call traffic also depends on the ASR and ALOC values. Plus, we must consider the performance during peak, off-peak hours and extreme traffic splashes (e.g. during Christmas, Easter, etc.). Thus, the total performance capacity must be greater than the maximum anticipated BHCA (busy hours call attempts).

Using the assumptions above (60% ASR and 3 minutes ALOC), we estimate that the number of minutes during peak hours is $140 * 60\% * 3 * 3600 = 907,200$ minutes.

The number of peak hours in a day highly depends on the customer types. For example, the number of peak hours for residential users is 2-4 hours while for business customers (e.g. call centers) it is 6 hours.

If to assume that there are 6 peak hours in a day, this results into $907,200 * 6 = 5,443,200$ minutes of peak traffic per day. Assuming that a month consists of 20 business days, we have $5,443,200 * 20 = 108,864,000$ minutes of monthly traffic.

Thus, to be on the safe side and allow for some reserve capacity, PortaBilling® is estimated to handle **100 million minutes** per month.



The results are calculated based on the **maximum** server's capacity, which is 140 CAPS. However, we strongly advise you to reserve some capacity of the billing server when you plan your infrastructure. This is required to handle traffic splashes and the traffic increase when your business grows.

SMS

Unlike calls, which last for some time, an SMS is a one-time event. Thus, a single Diameter request is generated for an SMS. PortaBilling can process about 200 SMS per second, which means that every second 200 users can send an SMS on your network. Assuming that 95% of all SMSs are successfully delivered, this results in **190 SMS per second**.

Internet sessions

An Internet session can last for a long period (e.g. for hours) without interruption. The service consumption is reported in session update (re-authorization) requests. Therefore, the number of authorization and re-authorization requests (transactions) received during the session is considered when measuring PortaBilling®'s performance.

The frequency of requests from P-GW to PortaBilling® depends on how quickly a user consumes the amount of data allocated for a session during authorization. The allocated data has a validity time; thus, P-GW also sends a re-authorization request after the validity time expires.

A single PortaBilling server with 32 active processing engines can process about 150 TPS (transactions per second). This means that every second 150 users' Internet sessions can be started or updated. The exact figures highly depend on your billing configuration (whether you provide capped or uncapped Internet service, whether you divide Internet traffic into rating groups), your overdraft protection settings (how many funds to lock for a session) and how quickly users consume the allocated data. Since it is not possible to know which service configuration you use, we can only **estimate** PortaBilling®'s processing speed, which has been calculated as **150 transactions per second**.

How many concurrent Internet sessions this translates to? This again depends on the frequency of requests sent by P-GW, the validity time and the percentage of simultaneously connected users. Given the processing speed of 150 TPS and assuming that 90% of all requests are successfully processed, this means that $150 \text{ TPS} * 90\% = 135$ user sessions are in a "connected" state. Depending on the validity time, the number of simultaneous sessions changes as shown in the table below:

Number of connected sessions	Validity time, s	Number of concurrent sessions
135	1800	$135 * 1800 = 243,000$
135	600	$135 * 600 = 81,000$

The results are calculated based on the maximum server's capacity, which is 150 TPS. However, we strongly advise you to reserve some capacity of the billing server when you plan your infrastructure. This is required to handle the traffic increase when your business grows.

Dedicated database server

In the recommended installation configuration for PortaBilling® the application servers are separated from the database for enhanced performance. The extra processing power and increased database in-memory cache on the dedicated server dramatically reduce the time required to execute many database queries.



Hardware considerations

The PortaBilling® billing engine uses multiple processors (or multiple CPU cores) in your server to process several requests in parallel, thus boosting performance – so adding another CPU or using quad-core CPU instead of dual-core provides a significant increase in performance. The amount of RAM is also very important, especially on database servers, while on a system with heavy load the disk subsystem becomes the component which determines the actual performance. Thus fast disk drives (15K rpm) and RAID controllers with adequate caching (for both read and write operations) are recommended.

High availability

PortaBilling® cluster

In addition to the ability to run several RADIUS processes in parallel on the same machine, you can also utilize multiple physical servers to process RADIUS requests. This provides extra performance, since incoming requests are distributed for processing among all the available servers. It also offers you improved reliability, because if one of the servers is down due to hardware failure, the remaining servers will continue operating. RADIUS cluster technology is included in the PortaBilling® with Oracle Real Application Cluster database and in the clustered version of PortaSwitch®

Technical details

Each server in the cluster runs two types of processes: dispatcher and processing engine. The dispatcher receives a RADIUS request from the network and sends it to one of the available processing engines. The actual business logic (e.g. checking whether it is a valid account/customer, calculating allowed session duration, or producing CDRs) is done by the processing engine. Normally there is a single dispatcher and multiple processing engines on each physical server.

One of the dispatchers is always active, while all the dispatchers constantly communicate and thus know which nodes in the cluster are alive. Should the server which runs the active dispatcher go down due to hardware failure, a new active dispatcher will be selected. The active dispatcher places requests into queues for the processing engines, based on their availability and current load. When the processing of a request is finished (meaning that xDRs are inserted in the database, balances are updated, etc.), the dispatcher receives confirmation from the processing engine and relays the response to the network node (this may be a simple accounting confirmation, or could include attributes such as “maximum authorized time”). If one of the servers in the cluster goes down, all

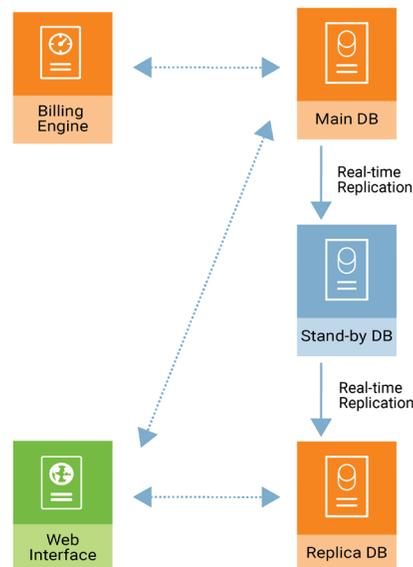
requests queued for processing engines on that server are processed by the remaining servers.

Performance

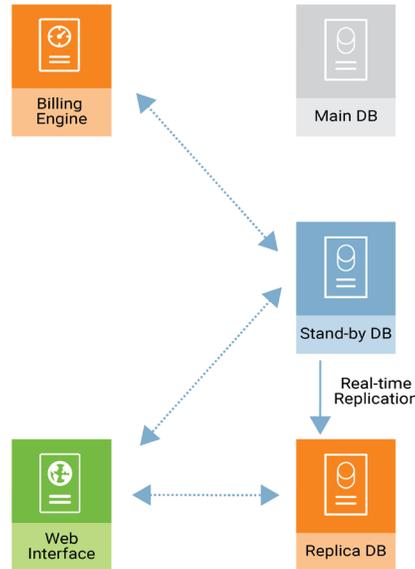
Since the load is distributed among all servers in the cluster, the total performance capacity grows linearly with adding more servers into the cluster. So if we add to a PortaBilling®, capable of **90** call attempts per second, another identical server and connect these two into a cluster – the performance of this two node cluster will be about **180** call attempts per second and by adding a third node it can be increased to **270** call attempts per second.

Using standby database for improved reliability

For quick disaster recovery in case your main database server goes down (e.g. due to motherboard failure), PortaBilling® can be equipped with a standby database server. This server is identical to the main database server, with the same software installed. The only difference is that during normal system operations the database runs in slave mode (with real-time replication), and so is fully synchronized with the main database.



When the main database is unavailable, the standby database can be switched to the role of main database and assume all real-time data access and modifications.



The standby server may be installed in the same co-location center (even in the same rack) as your main database server. This provides easy administration and, in case of a switch from main to standby database, the standby server is able to simply take over the main server's IP address. In such a case, the change is transparent for all elements of your VoIP network.

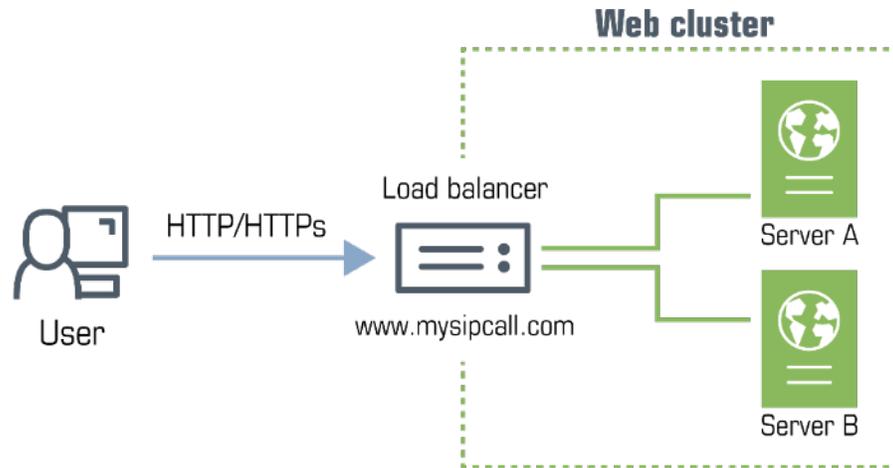
However, should your co-location provider be having a problem with network connectivity, both servers will become inaccessible to your VoIP network installed elsewhere. This can be avoided by installing the standby database server and additional billing server in a different physical location, connected to a different network (and preferably to a different ISP). This is so that even if the whole hosting center including your main database server is down, you can still continue network operations via the servers at the remote location, as they will not be affected by the outage.

Web cluster

The purpose of a web cluster is to allow the uninterrupted processing of web requests from administrators, end users or external applications in case of hardware failure on a physical server. It also enables load sharing among multiple servers, which allows the system to easily scale up to handle a higher total number of users who access the system via a web interface or external applications, or who access data via XML API.

For outside users (or applications), the cluster is accessible via a single virtual IP address. For instance, if the cluster's virtual IP is 1.2.3.4 and one or more domain names such as `www.mycall.com` are set to point to that IP, when a user types "`www.mysipcall.com`," his request is delivered to the cluster and then processed by one of the nodes there. The address to

which all users and / or applications connect is always the same, so if any server is down, no reconfiguration is required in external applications or via the end-user's web browser.



All incoming sessions are distributed among available nodes by the load balancing application according to a predefined algorithm. The default algorithm is the “round robin,” according to which requests are sent to each of the instances in turn.

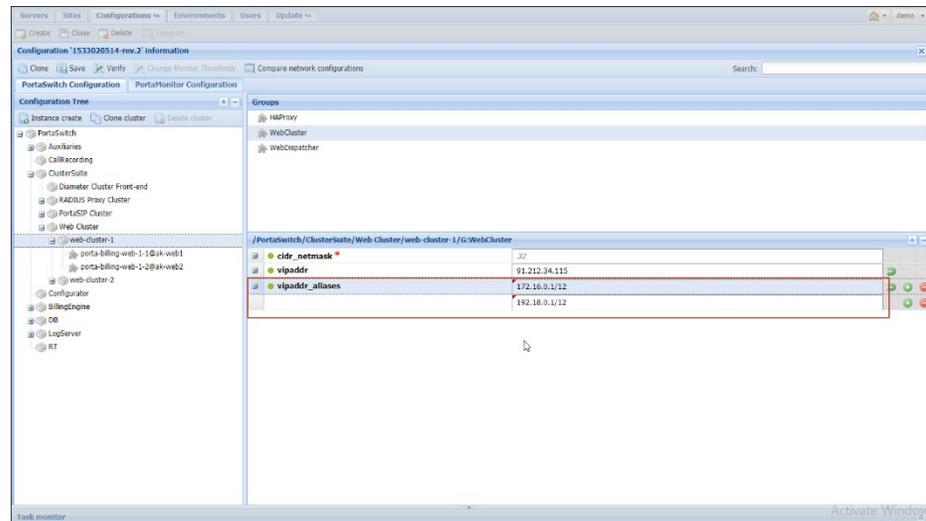
The load balancer also monitors the state of web instances within the cluster. Whenever an instance is unavailable, the load balancer adjusts the request distribution. The load balancer is deployed on every web instance; however, it is only active on the instance with the virtual IP address.

The cluster starts with a minimum number of two servers. The number of web servers that can be added to the cluster is unlimited.

IP aliasing

SSL encryption is among the key requirements for websites. Encrypted sites provide a secure connection between the web server and the user, ensure the safety of user sensitive information (e.g. credentials, credit card data, etc.) and have higher rankings among search engines. In order to maintain a secure website over SSL encryption, a dedicated IP address is required.

Web cluster supports multiple IP addresses as aliases. This enables you to allocate dedicated IP addresses for your environment owners and resellers to build separate, encrypted websites.



With a dedicated IP address, your customers and resellers can:

- Create websites under their own domain names;
- Secure the user data transmission by using their own SSL certificates; and
- Protect their websites from being blacklisted by managing the firewall configuration.

Thereby you organize white label operators to work independently. You also prevent network bottlenecks from occurring when a mass of traffic arrives at the same host and port.

2. System concepts

Overview

Converged billing

PortaBilling® supports multiple services and service types. This means that as different types of services (e.g. voice calls, WiFi connectivity and messaging) are provided to your users, PortaBilling® collects data about all of them, processing and rating it according to the billing configuration. It then provides your customers with a consolidated bill (and your administrators with a unified customer management interface).

Billing events

The main unit of billing information is a billing event – a notification that a service has been provided to a customer in the outside world, and that this customer should thus be charged for that service.

For many services (such as SMS messaging), a single billing event is represented by a single notification message to billing, while for others information about one event (e.g. a completed phone call) is split into multiple notifications from different network elements.

Four entities involved in the billing

Four separate entities will be billed for every event registered in the system:

- Account (identified entity using the service).
- Customer (an account owner).
- Reseller (optional).
- Vendors (service providers, e.g. your termination partners or service providers for incoming telephony lines, both local and toll-free).

This allows you to:

- Bill the end user (the owner of a prepaid calling card, phone number or office VoIP gateway).
- Bill the customer to whom the end user belongs.
- Bill the reseller who owns that customer (if applicable).
- Calculate your expenses for termination of this call to the vendor.

xDRs

CDRs (Call Detail Records) were the real foundation of any legacy billing system, which had to receive, store and analyze them. Today, when most

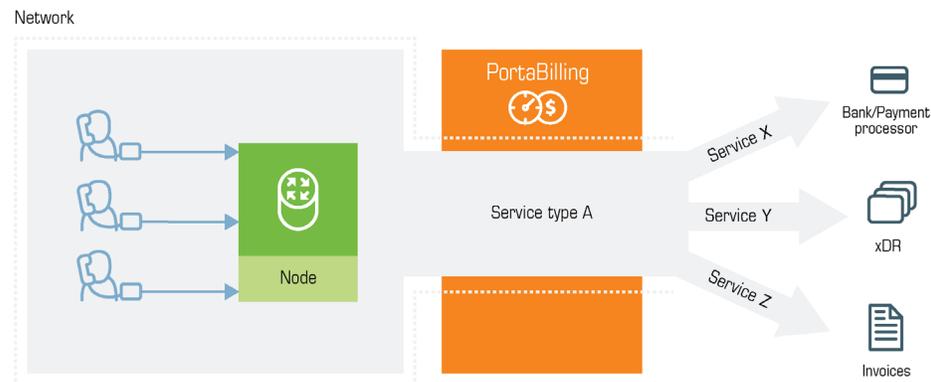
service providers are expanding their service portfolio, the CDR concept is either becoming inapplicable to some of their services (e.g. selling pizza over the Internet), or it no longer fits their business model (e.g. providing bulk free calling minutes). PortaBilling® supports this by introducing the concept of xDRs – eXtensible Detail Records, which can store information about various types of billing events. CDRs are, therefore, only a subset of all possible xDR instances. The same applies to EDR (Event Detail Records), data which describe a single event such as an SMS message; they too are just a subset of xDR data, and can thus be easily processed by PortaBilling®.

Service types and services

PortaBilling® enables you to offer multiple services to your customers, so you can broaden your range of services while providing clear transaction and invoice information to your subscribers and administrators.

There are two entities involved in the billing configuration:

- Service types
- Services



A **service type** is a description of the physical service provided to an end user. A service type primarily defines the way in which the billing engine interacts with network equipment to ensure that the service is authorized, provided and rated correctly. VoIP calls or WiFi Internet access are examples of service types. The list of available service types is determined by the billing engine capabilities in a given release, as new service types appear when proper support is added to the billing engine internal modules.

A **service** is a description of the user's activities from a business perspective. Services are the foundation of your billing configuration, and new services may be defined according to your business model. This permits flexible branding of your services to customers and opens the possibility of bundling various services into one product. You can choose names and billing parameters for new services as desired. For instance, the same physical service, such as Internet access, can be sold to one category

of users as “High-speed Internet”, and to another category of users at a different rate as “Broadband connectivity”. Moreover, the **same** physical service can be sold to the **same** customer as multiple entities: for example, the use of WiFi in the customer’s office and in an airport are treated as two separate services with a different price. A breakdown of charges per service is provided on the invoice and in the xDR browser.

Rating base

When information is received about the use of a particular service type, the billing engine must process it and decide which attribute or set of attributes are to be used later to calculate the charged amount. This may be quite different depending on the service type: session duration, amount of transferred bytes, counter values, and so on. The logic determining which attributes are used for billing is called the rating base. For instance, the rating base for the **Internet Access** service (when charging for a data transfer) will define what exactly constitutes the “traffic” which the customer is being charged for: the amount of downloaded data, the amount of uploaded data, or both.

The most commonly used rating bases are:

- **session-time** – These include phone calls, Internet access sessions and other services where the basis for calculating charges is the duration of time in which a customer used the service.
- **quantity** – These services include data transfer, text messaging, pay-per-view movies, or pizza delivery orders. The duration of the service session is either unsuitable or irrelevant for billing purposes. Rating is done based on the type of service element used (e.g. if the movie purchased was a “New Release” or “Classic”) and the quantity (did you order two pizzas, or just one?).

Each service type supports at least one rating base. Potentially, however, it can have more than one, e.g. the **Internet Access** service can be charged based on either session time or the amount of data transferred.

Every new service has a name and several important parameters:

- The service type it uses (see below for a description of available service types).
- The rating base, defining which set of attributes will be used to calculate the charged amount.
- A base unit name; base units are those units in which the amount of the service “used” is reported to billing (e.g. information about session duration is usually given in seconds, while information about the amount of data transferred is given in bytes). You cannot change the type of base units, since this is typically defined by the configuration and capabilities of your network equipment.

You can, however, give it a name which is easy to understand, or even put it in your user's local language.

- Billing units are used for all functions in PortaBilling®. These are the units in which your pricing is specified, in which customers see the amounts on their invoices, and so on. For instance, the billing units for most telephony calls are minutes, while in the case of data transfer services it could be kilobytes (for a GPRS data service) and megabytes or even gigabytes (for a broadband Internet connection service).
- Obviously, you need to define how your billing units relate to your base units. PortaBilling® gives you full flexibility to decide whether 1 kilobyte will equal 1,000 or 1,024 bytes for billing purposes. Rating schemes for other services may be defined with the same degree of flexibility.

Measurement, base and billing units

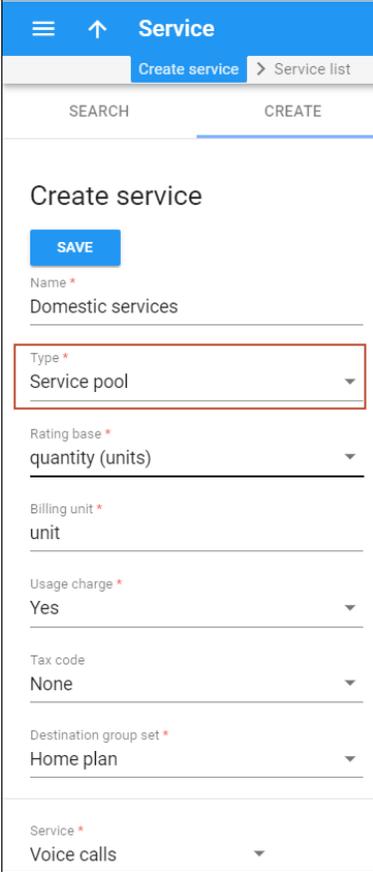
Measurement units are the units in which the billing engine receives information from the network. Obviously, the billing engine needs this information in order to interpret amounts correctly. However, a PortaBilling administrator rarely comes into contact with these units, since for the purposes of adjusting the rating configuration he would use the following:

- **Base unit** – The smallest elements that can be rated. This can be thought of as the “rounding threshold”. For example, if a customer has downloaded 758 bytes, but the base unit chosen is a kilobyte, for all further rating calculations he is considered to have downloaded 1 kilobyte.
- **Billing unit** – The unit in which it is most convenient to express the amount of services used. The price of the service will also be defined using this unit.

The base unit must have either the same or less precision than the measurement unit, and the billing unit must have either the same or less precision than the base unit. Thus while in some cases the units will be equal (e.g. call duration is reported in seconds and the base unit for the voice calls service is also one second), in other cases you can scale units to make them more convenient for your administrators and customers. Broadband Internet services are one good example. Although network equipment always reports the amount of transferred data in bytes, it would certainly be very confusing to report the total downloaded data as “234564474 bytes”, or inform customers that “your current rate is \$0.000000014 per byte”. They will be much happier to see “228 megabytes” or “\$1.50 per gigabyte”. So in this case you would choose a megabyte as the base unit and a gigabyte as the billing unit.

Service pools

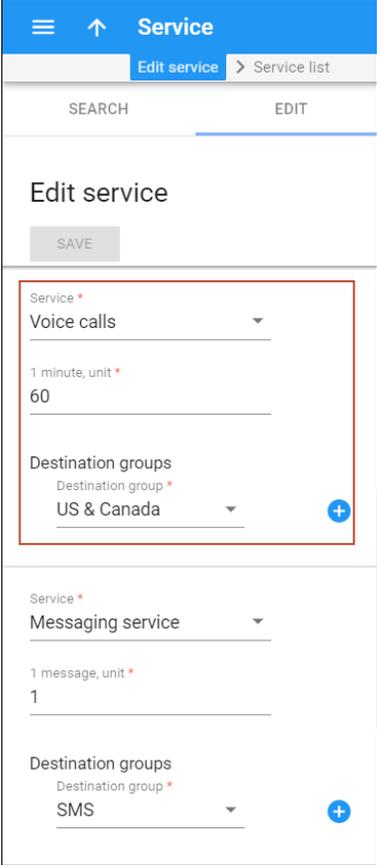
To provide several services to your customers within a single quota, consolidate them in service pools. However, every service within the pool can be measured differently (e.g. voice calls are measured in minutes, Internet access – in megabytes, etc.). Therefore, service pools are measured in unified *virtual units*.



The screenshot shows a mobile application interface for creating a service. The top bar is blue with a menu icon, an up arrow, and the text 'Service'. Below the bar are two tabs: 'Create service' (active) and 'Service list'. The main content area is titled 'Create service' and has a 'SEARCH' button on the left and a 'CREATE' button on the right. A 'SAVE' button is located below the title. The form contains the following fields:

- Name *
Domestic services
- Type *
Service pool
- Rating base *
quantity (units)
- Billing unit *
unit
- Usage charge *
Yes
- Tax code
None
- Destination group set *
Home plan
- Service *
Voice calls

Virtual units are used to express the cost for using a particular service that is included in a pool (e.g. a minute's worth of calls costs 5 units while an SMS message costs 1 unit). If your pricing changes, you can conveniently re-configure it within a service pool.



The screenshot displays the 'Edit service' interface in a mobile application. At the top, there is a blue header with a menu icon, an up arrow, and the title 'Service'. Below the header, there are two tabs: 'Edit service' (active) and 'Service list'. The main content area is divided into two sections. The first section is titled 'Edit service' and contains a 'SAVE' button. Below this, there are two service entries. The first entry is 'Voice calls' with a unit of '1 minute' and a value of '60'. The second entry is 'Messaging service' with a unit of '1 message' and a value of '1'. Both entries have a 'Destination groups' section with a dropdown menu and a plus sign button.

Additionally, you can fine-tune the pricing for the same service within the pool. For example, make calls to Europe cost 10 units per minute while calls to Germany will only cost 8 units per minute. To do this, select the same service for different destination groups and then configure the pricing.

Service pools can only be used in volume discount plans and be applied for quotas and service wallets. Please refer to the [Quota / Balance Allocation for Multiple Services](#) section for details.

To summarize, service pools have the following parameters:

1. Service pools apply for a complete destination group set.
2. Each service within a pool is defined for a destination group.
3. Service pools are measured in virtual units. Virtual units are used to express the cost for using a particular service and are mapped to this service's units (e.g. a minute's worth of calls costs 5 units while an SMS message costs 1 unit). Note that the number of virtual units in the mapping must be higher than the number of service units.
4. To fine-tune the pricing, you can add the same service multiple times for different destination groups and define a different unit mapping. Note that the service-destination group pairs within the service pool must be unique.

5. You can create a single multi-service wallet or quota within the same volume discount plan.

Supported service types

Once installed, PortaBilling® supports the service types shown in the table below. The Rating Base column refers to the applicable rating base options (S means “session-based” and Q means “quantity-based”):

Name	Rating Base	Description
Conferencing	S	Rating conference calls via a PortaSIP® Media Server (or some conferencing server).
Data Service	Q	Rating data transfers using the amount of data transferred as the billing parameter.
Dial-up Internet	S	Dial-up Internet access sessions, rated based on session duration.
IPTV	Q	IPTV services, like pay-per-view movies, rated based on the number of views.
Internet Access	S, Q	Internet access sessions (DSL, PPPoE, etc.), rated based on session duration or the amount of transferred data.
Measured Service	Q	Rating the amount of an allocated resource (active calls, IP Centrex phone lines and other), charges are based on the number of resource units consumed.
Messaging Service	Q	Rating messages (text, SMS, MMS, other); charges are based on the number of messages sent.
Mobile Network Provisioning	Q	This enables subscriber data provisioning to mobile networks. When enabled for an account’s product, the External System Provisioning Framework (ESPF) provisions both policy / service quality parameters and account management operations accomplished in PortaBilling® to the mobile carrier’s core.
Quantity-Based	Q	Generic quantity-based service type; can be used to apply charges for any service use expressible in numeric form (e.g. number of pizzas ordered).
Session-Based	S	Generic time-based service type; can be used to apply charges for any service use based on the length of time the service was accessed.

Voice Calls	S	Rating telephony calls (incoming or outgoing) made via a PortaSIP® Switching Server, VoIP gateways or other equipment.
Wi-Fi	S	Wireless Internet access sessions, rated based on session duration.
Service Pool		The set of services provided within a single usage plan.

Service provisioning

PortaBilling® is a unified place to configure and manage the services you provide to customers.

All services can be easily configured and assigned to your customers via the administrative web interface. It represents a convenient tree-structure view of available services and features (e.g., “Call Waiting” or “Maximum Allowed Bandwidth”), so you can manage them and quickly see which of them are enabled / disabled for a particular product, customer, etc.

ACCOUNT SERVICE CONFIGURATION

12065551234

SERVICE: INTERNET ACCESS

SERVICE: VOICE CALLS

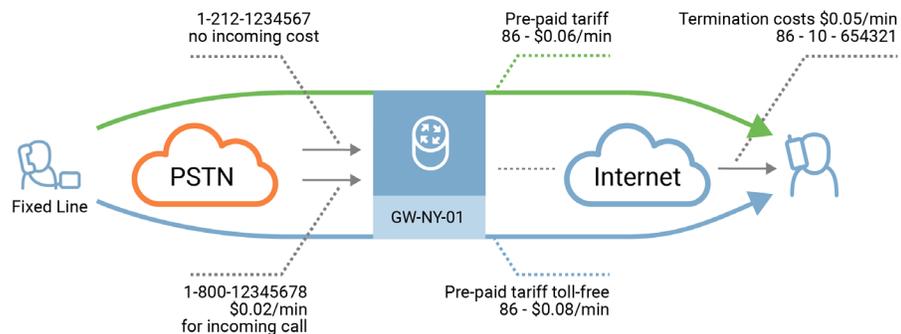
- INCOMING CALLS
 - CALL WAITING
 - CNAM CALLER INFO
 - FOLLOW-ME
- OUTGOING CALLS
 - OVERRIDE IDENTITY
 - CLI = 12065559000

PortaBilling serves as a single interface for the service configuration: when changes to the service configuration are made through web GUI, they are automatically active and used by PortaBilling®. If there are external components without real-time link to PortaBilling® (e.g., HLR), they are updated via the provisioning framework almost immediately. There is no need for manual provisioning using additional applications or interfaces. This allows instantaneous service activation and drastically reduces the administration overhead.

Products

A product is a combination of services that you provide to a customer for a price. For instance, the same product may provide both wireless Internet access and cheap international phone calls from an IP phone. Or perhaps you will decide to sell calling cards with 6 cents/minute calls to China via a local access number in New York, and 8 cents/minute calls to China via a toll-free line. In this case, your product will include two types of service:

- access via the local New York number, and
 - access via the toll-free line,
- with price parameters associated with each service.



The rating table (available on the **Usage charges** panel) is the main component of a product definition. It specifies which services your customers have access to, where on your network they may use these services, and how they should be charged for them. It allows you to specify the following parameters which define a service consumption point:

- The type of service provided.
- The node on which the service is used. What exactly does “node” mean in this context? If, for example, a customer calls to gateway A, enters his PIN, and makes an outgoing call which is terminated on gateway B, is he using a service on node A, node B, or both? The correct answer is that the service is regarded as having been provided **at the point where authorization was performed**. In this example, since PIN authorization is performed on node A, it is node A which must be listed in the rating table.
- Identification of the access code (method) on that node. This parameter allows you to use different rate plans for the same service. For example, you may choose a rate plan according to the PSTN access number (local or toll-free) that the customer has dialed. Or you may use different rate plans for outgoing, incoming and forwarded calls in your SIP calls service using the OUTGOING, INCOMING and FOLLOWME access codes, respectively. (While for services such as prepaid cards the access

code is a number, for other services any string may be used, so long as it is one provided by the application handling the call).

- Originating line information (this is applicable only to the voice call service). You can separate rating entries based on originating line information (e.g. whether the call was made from a home phone or a pay phone). Make sure your telecom provider supplies you with this information in the call setup. For more information please refer to the **OLI-Based Billing** section of this guide.

Below is an overview of the product bundling and rating possibilities offered by PortaBilling®. The MegaStar product allows customers to make and receive calls from their IP phones, connect to a WiFi network and be charged for time spent online, and exchange instant messages. You may assign a different rate plan for each service, as shown.

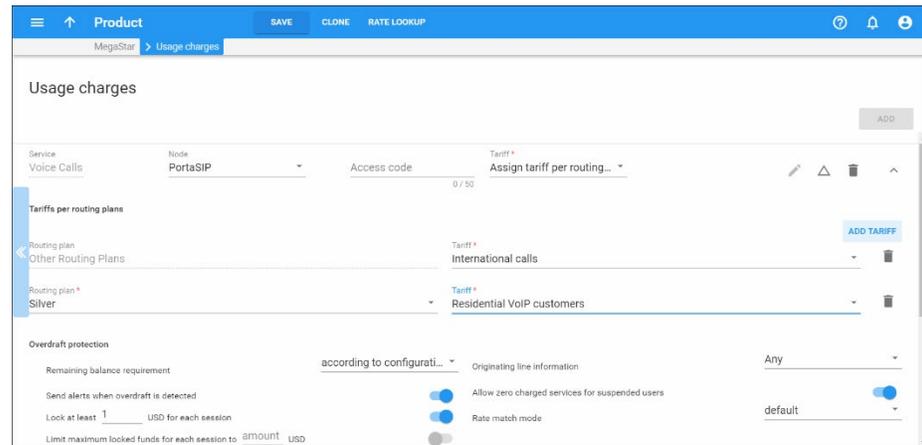
Service	Node	Access code	Tariff
Wi-Fi	Mikrotik WiFi	Any code	Wi-Fi service
Messaging Service	PortaSIP	Any code	IM Services
Voice Calls	PortaSIP	INCOMING	IP Centrex - incoming
Voice Calls	PortaSIP	Any code	IP Centrex - outgoing

Applying tariffs based on routing plans

Normally, the customer is charged the price which is pre-set in his price list. Naturally, this price is calculated based on estimated average termination costs, but does not directly depend on which carrier is used for a specific call. Since you may have different termination vendors with various prices and levels of audio quality, how can you give your customers the freedom to choose whether to make phone calls via low-cost carriers (and thus pay low rates), or use high-quality carriers (and so pay premium rates)?

Routing plan selection is available for services based on the **Voice Calls** service type. It allows you to select a tariff for charging a customer based not only on the “incoming” point to your network, but also on the “outgoing” entry, i.e. the type of route used to terminate a call. Each type of routing (low-cost, premium, etc.) is defined as a separate routing plan, and multiple routing plans can now be associated with a single rating entry in the product definition. A specific tariff will be assigned to each routing plan, and this tariff will be used when a customer

subscribing to a specific plan places a call. For more details, see the *Routing Plan Selection* chapter.



You can still use the “single tariff per rating entry” mode. Since this provides faster performance, it is a good idea to use it when the real-time selection of routing plans is not required. In this case, you can still assign a default routing plan to an account, and this will be used when the end user makes a phone call.

“Single tariff” mode is active by default when you create a new rating entry. If you wish to apply a specific tariff based on the routing plan used, you have to manually select the “Assign Tariff per Routing Plan” option in **Tariff** list.

Add-on products

General concept of add-on products

Add-on products are a great way to sell additional features / services to your customers and to boost your sales revenue.

Add-on products come with a feature and a subscription attached to them. You only need to assign the add-on product to an account. The feature activates and the customer is automatically charged for it.

Add-on products have the following features:

- Multiple add-on products can be added to an account, creating a desired service configuration.
- Some service features may be defined, so the actual service configuration of an account may have a combination of the service features / attributes of all the products (main and add-ons) and service features / attributes configured on the account directly.

- They may include a definition of extra entries for volume discount plans to be applied to add-on products.
- They often have subscriptions included – so, an account that consumes this add-on product will obligatorily have those subscriptions assigned.

Once a single process for changing a service configuration and applying charges for it is established, there will be an opportunity to:

- Efficiently manage a set of products.
- Avoid revenue leakage (when a service is activated, but a customer is not charged for it).
- Avoid billing disputes (when a customer is charged for a service that was never provisioned).

Add-on products service configuration

Correct service configuration is an important step of product setup since it:

- Ensures that customers are supplied with the quality services that they need.
- Decreases the workload on administrators as it is easier to manage those products in future.

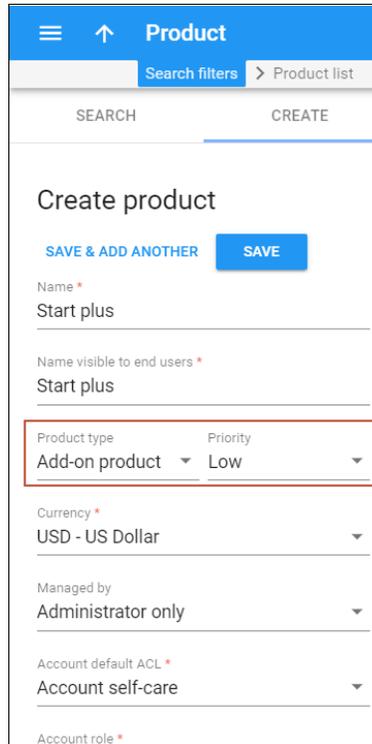
The configuration itself starts from the product, which consists of services consumed by end users. Each service is defined by service type (for example, Voice Calls), which describes the physical service provided to an end user. Every service type consists of service features such as **Call waiting**, **CNAM**, etc., which can be activated or deactivated by an administrator. Apart from enabling / disabling for certain features – when they are enabled, additional parameters called “service attributes” must be defined. For example, when the **Call via IVR** feature is enabled, that’s not enough for that functionality to work properly. Specific IVR applications must also be selected, and these will be used for calls by the end user. These applications should be selected from the **Voice Application** list on the **Outgoing calls configuration** panel of the product. In this case, the **Voice Application** is a service attribute of the **Call via IVR** functionality.

When configuring services, the services that your customers can use and how they use them can be specified. The service configuration is calculated based on the settings performed within assigned products (main ones and add-on ones).

When configuring add-on products, keep in mind the following peculiarities of the add-on product functionality:

- **Priority.** Add-on products have priority over main products, so no matter what is defined within a main product it will be overridden by settings defined within the add-on products (this

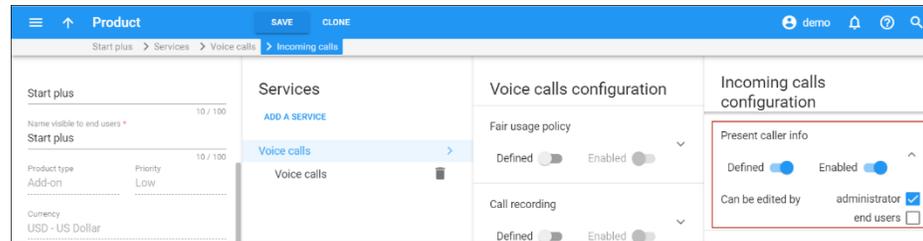
applies only to the options that are supplied with add-on products). In order to differentiate between the add-on products there is also an “override priority” parameter for each add-on product. If more than one add-on product is assigned to an account, they will be sorted according to the specified priority.



Product type	Priority
Add-on product	Low

For example, an account with a “Basic Call” main product is assigned a “Smart Call” (low priority level) add-on product with the **Conferencing** feature disabled. Then the customer decides that he needs the conferencing service, so he orders the “Conferencing+” (high priority level) add-on product with the **Conferencing** feature enabled. In order for the system to apply the correct values, it weighs the priority level of each add-on product. As a result of service calculation, the conferencing service is enabled for that account.

- **Service Features.** Each add-on product provides only those functionalities that are explicitly defined within its service features. If you need to specify the product default value of a certain feature, move the slider next to the service feature name.



All other service features that are not defined within the add-on product are ignored during the service calculation. In this case, the definition of the feature is taken from another product (in which it is explicitly defined) or from account settings. On the web interface, such service features can be identified by an empty check box on the left of the service feature name (this is the default position of the check box).

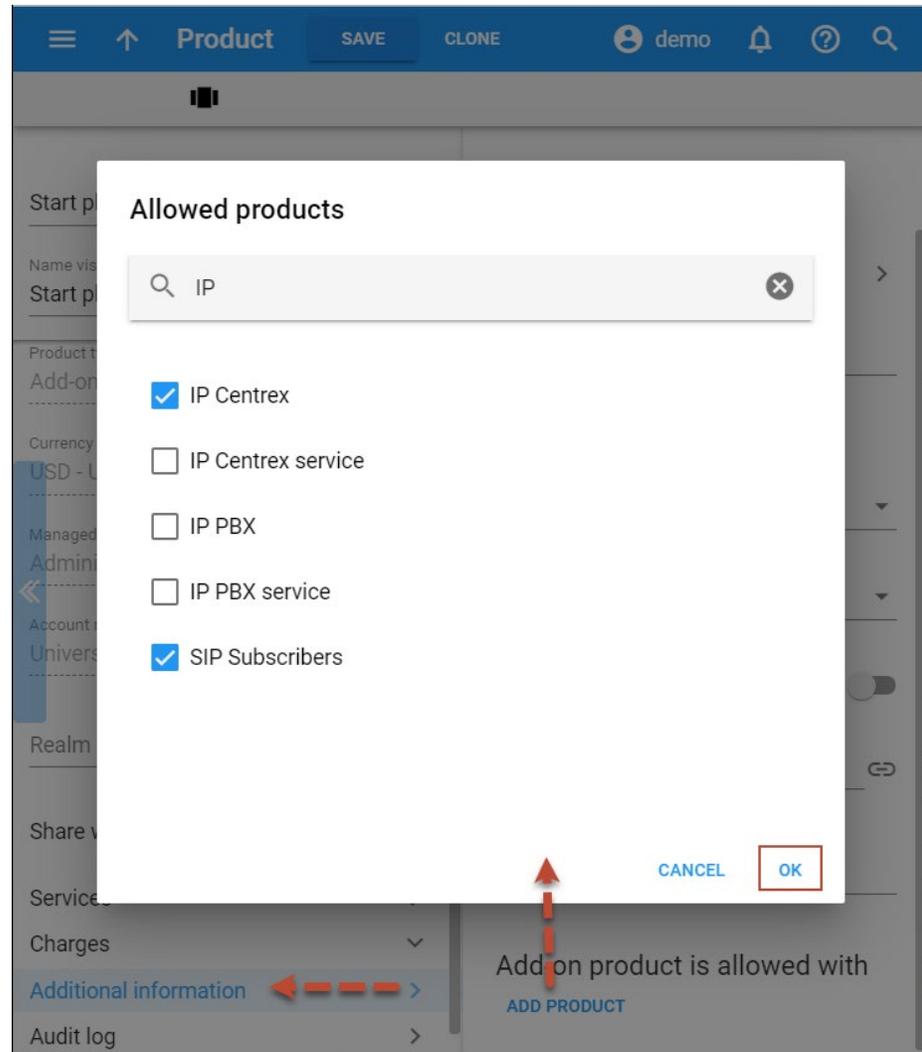
- **Service Attributes.** All attributes can be sorted into three categories:
 - Those that are defined on the product level only – they are usually the ones that have some business value and should be controlled by a product.
 - Those that are defined for every account separately, e.g. the list of follow-me numbers or the value of “Outgoing CLI” should be defined for each account individually.
 - Those that can be defined on both product and account levels.

For example, the maximum number of forwarded calls is normally defined on the product level, so residential customers have it set to 3, and IP Centrex customers to 10. However, for some accounts, it may be set to a higher value (for example, there are 12 people in the customer’s sales department and the customer wants to forward calls to all of them).

Note that a service feature that is defined on the account level overrides the one provided by the product.

- **Product compatibility.** Every add-on product has a list of allowed main products with which this add-on product is compatible and with which it can be assigned to an account.

For example, the main product “EasyCall” only has basic and inexpensive services. Of course, every customer can order additional features, but to do this, the main product must be upgraded to an advanced version – i.e. the “Smart Call” product, with which additional services “a la carte” can easily be assigned. In this case, all add-on products will have the “Smart Call” product allowed.

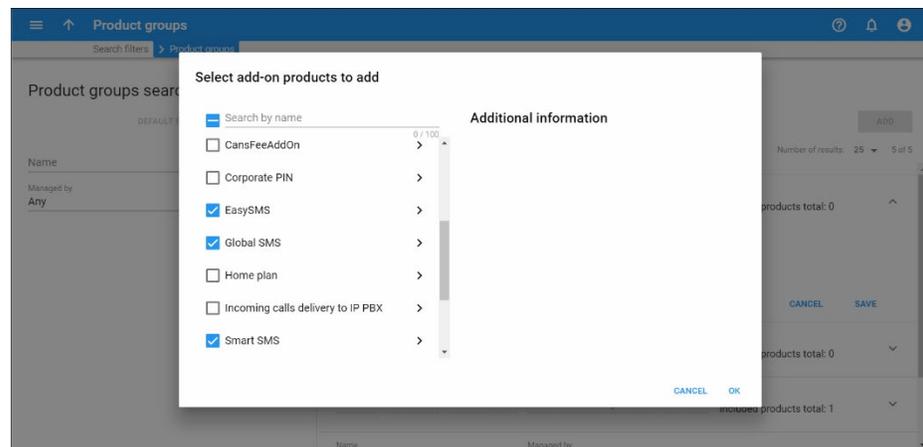
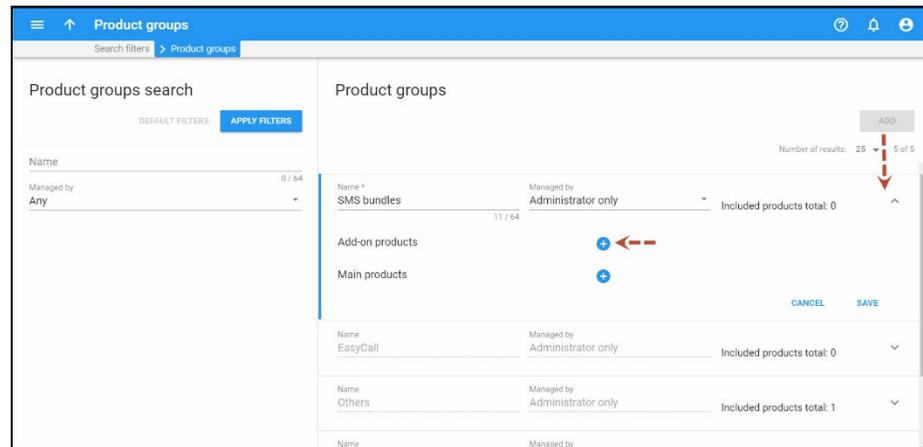


- **Service features conflict.** A situation where the same service feature is set by different add-on products with the same priority level cannot occur. The system does not allow such configurations to be saved and will notify an administrator concerning which features and add-on products caused the conflict. For example, an administrator tries to apply the “Unconditional forwarding” add-on product with **Call Waiting** functionality disabled and the “Advanced forwarding” add-on product which has **Call Waiting** functionality enabled. Both add-on products have the same priority level, so saving this configuration will not be allowed since it is unclear whether the **Call Waiting** functionality should be enabled or disabled.

Product groups

When you provide various add-on products and allow end users to sign up for add-ons via the account self-care portal, you may also want to provide them with effective product management tools.

This can be done by organizing your add-on products in groups. Now administrators or resellers can create product groups and assign corresponding add-ons to them via the admin interface / reseller’s self-care.



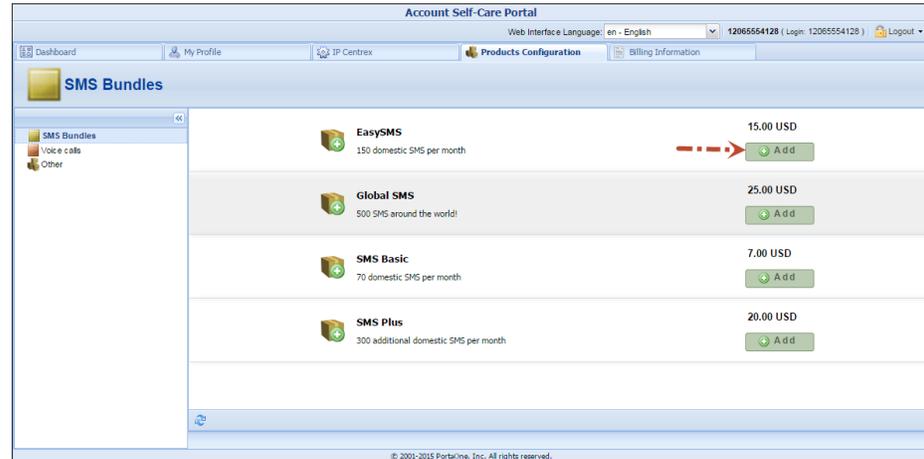
Then these groups are available to end users on the **Products Configuration** tab of the account self-care interface. A user can easily find a desired add-on by entering the corresponding group that the add-on belongs to.

Consider the following example:

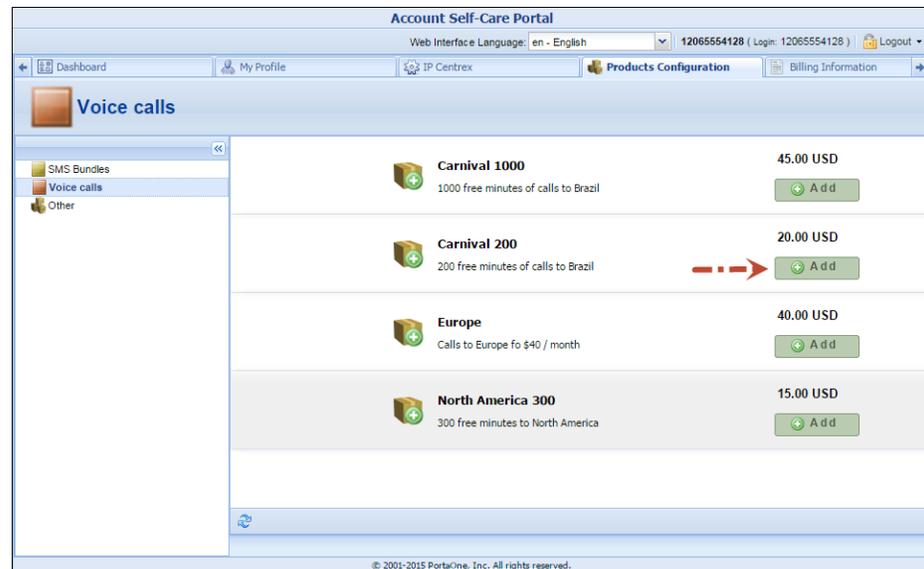
Customer John Doe has 50 free domestic SMS and 100 free minutes of calls to US and Canada destinations included in his main product. He

wants to increase the number of domestic SMS to 200 and add 200 free minutes of calls to Brazil.

He visits his account self-care interface and goes to the **Product Configuration** tab. He selects the SMS Bundles group, finds the “Easy SMS” add-on product that provides 150 domestic SMS messages and signs up for it.



Then he selects the **Voice calls** group and signs up for the “Carnival 200” add-on product that provides 200 free minutes of calls to Brazil.



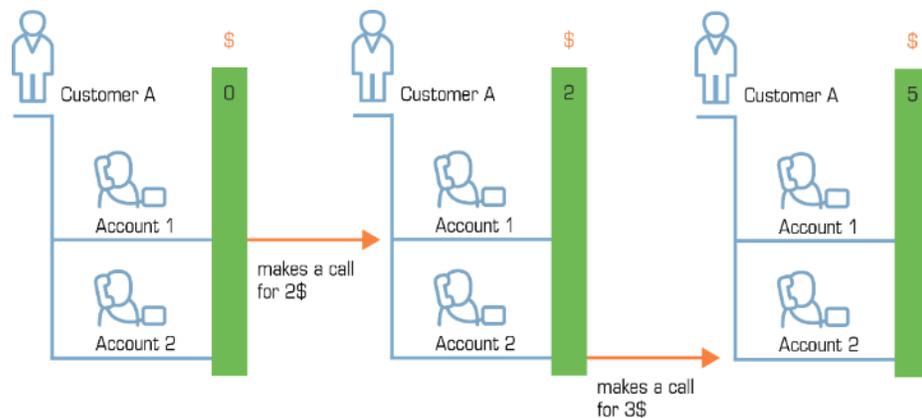
Now John Doe has his desired product configuration and can further enjoy its services.

Thus, product groups enable end users to perform a quick and easy product search and self signup, thereby decreasing the load for your administrators and customer care staff.

Retail customers

A customer is an entity (an individual or a company) who is responsible for using the service via one or more accounts. For instance, when a family signs up for residential VoIP services, a man, his wife and his brother will all have individual phone lines, each represented by an account. The person who signs the contract (and who is responsible for paying the bills) will be entered as the customer, and all the accounts will be listed under him.

When a credit account is charged for using the service, this directly affects the customer's balance, and the activities of all accounts are included in the consolidated invoice. This allows quick tracking of service use for a group of accounts, or the application of a credit limit for the whole group.



NOTE: If a call was made by a debit (prepaid) account, this does not affect the retail customer's balance (see below), since prepaid accounts (cards) are normally paid for by your customer at the time of purchase.

Typically used for: business end-customers (when more than one account per customer is required)

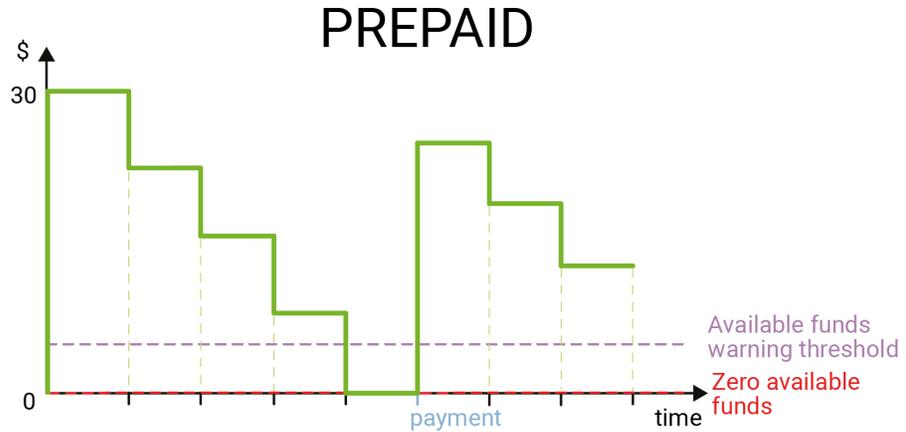
Customer balance model

Customers are divided into prepaid and postpaid categories, depending on how their balance is controlled.

Prepaid customers

For customers who have a prepaid balance model, they first pay for the services and then use them (e.g. customers will not be able to make outgoing calls or surf the Internet without topping-up the balance first).

The balance for prepaid customers is shown as a positive value and indicates how much of the services can be used with the funds currently available.



While consuming the service, the amount of funds decreases. When it reaches zero value, no more services can be used.

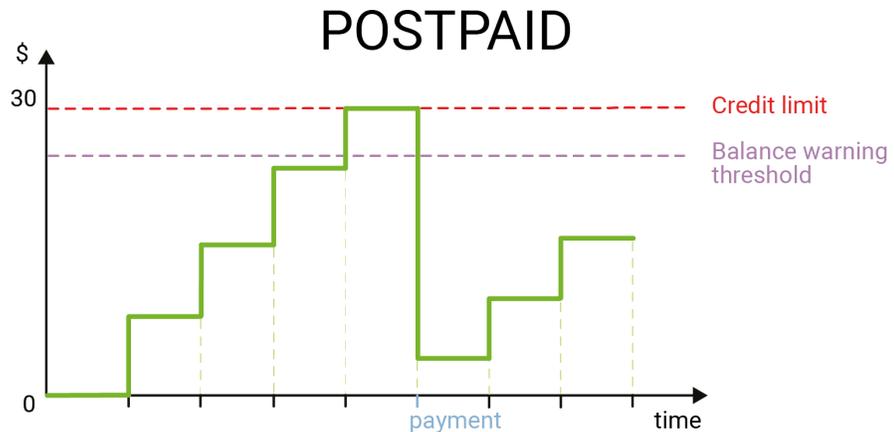
Once a payment is made, the amount of available funds increases.

Typically used for: calling card services.

Postpaid customers

For customers who have a postpaid balance model, services are provided to them on credit. This means they use the services and then pay for them on the basis of invoices generated at the end of a billing period.

The balance for postpaid customers indicates the amount of money owed to the service provider. Once the invoice is paid, the balance decreases.



To prevent misuse of services and fraud traffic (e.g. the account’s credentials were stolen and lots of traffic was sent to the network), it is highly recommended that you set up a credit limit for postpaid customers. When the customer’s balance reaches that credit limit, this customer’s accounts will no longer be allowed to use the service.

Typically used for: businesses (e.g. companies having several phone lines and paying consolidated invoices for the total amount of services used).

Please find the usage of customer balance models different business scenarios in the table below:

Description	Customer Balance Model	Account Type	Business Scenario Examples
The customer receives and pays consolidated bills for services used by all customers’ users. The customer balance and credit limit are shared among all users under that customer.	Postpaid	Credit subordinate *	IP Centrex, messaging service, triple-play service provisioning for residential customers
The customer prepays for service usage and doesn’t receive bills (though can receive statements). The customer’s available funds are shared among all users under that customer.	Prepaid	Credit subordinate *	Residential VoIP, messaging service, wholesale termination, IPTV and Internet access provisioning
One customer can have multiple individual users. Each user prepays for their service usage. No bills or statements are issued. One PIN / phone line with individual balance per user.	Any (doesn’t matter for debit accounts)	Debit	Prepaid calling cards, PINless dialing, callback, Hotspot and Viber-like services. Also can be used for social networks users

* It is possible to define individual credit limits for particular users under the customer (in this case users top up their accounts themselves with a credit card, either manually or automatically.) This account’s service usage

affects both the account's balance and the customer's balance or available funds.

Customer hierarchy

Organizing the customer hierarchy, administrators can easily configure the distributed infrastructure of a company with several functional departments (for example, the support and sales departments) so that their billing-related activities are maintained separately. Moreover, features such as extensions and huntgroups have only to be configured once before they can be shared among all of a company's departments.

This is achieved by introducing two types of customers within a hierarchy:

- **Main Office (HQ)** – This defines the “main” customer in the group for which the basic service configuration is done. All extensions and huntgroups added for this customer become available for all of its linked customers.
- **Branch Office (site)** – This defines the “subordinate” customer created under the Main Office (HQ) customer. This customer inherits all of the main customer's extensions and huntgroups.

Customer hierarchy is also useful for when an administrator needs to configure a company's geographically-distributed infrastructure. In this case, the company has multiple independent offices and requires individual billing statistics for each of them.

Consider the following, for example. The company has its main office on shore and its petroleum towers at sea. The personnel who work on the towers need to communicate with the main office via extensions and they also make calls to the main office's huntgroups. Their chief accountant wants each tower's billing and payments activities managed individually.

The administrator creates a Main Office (HQ) customer in PortaBilling and configures the extensions and huntgroups for it. Then he creates a group of Branch Office (site) customers and links them to the Main Office (HQ) customer created earlier. Since the extensions and huntgroups defined for the Main Office (HQ) customer are available for this group of linked customers, no additional configuration is required.

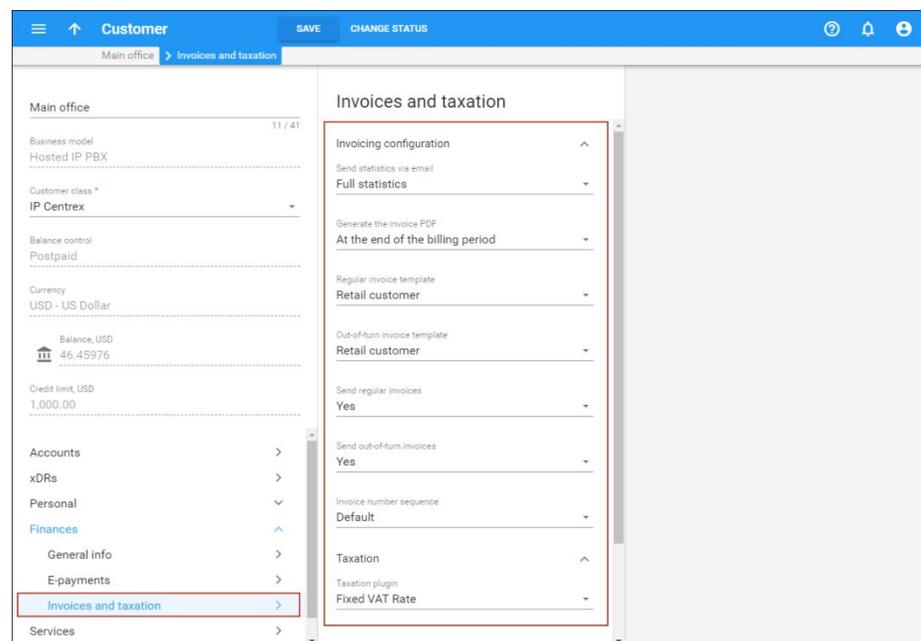
The administrator then configures the billing-related activities (such as invoice generation, balance management, amount due tracking) for each customer he created, and billing-related statistics for each of them can be tracked individually.

NOTE: To preserve individual invoice generation for linked customers, set the **Include New Branches Into Consolidated Invoice** option to **No** on the Configuration server.

Consolidated invoice for customer hierarchy

Let's say a company has a main office and a group of branch offices. All of the offices share the same IP Centrex environment (e.g. personnel can communicate across all of the offices via extensions, or make calls to other offices' huntgroups). The company owner wants to pay the communication charges for all their offices in a single bill.

An administrator creates a Main Office (HQ) customer and a group of Branch Office (site) customers, and links them together. Then the administrator enables both statistics and invoice generation, and defines the taxation settings for the Main Office (HQ) customer. The billing period must be the same for all linked customers in the group.



To avoid discrepancies during tax calculation, disable tax calculation for all Branch Office (site) customers within the group.

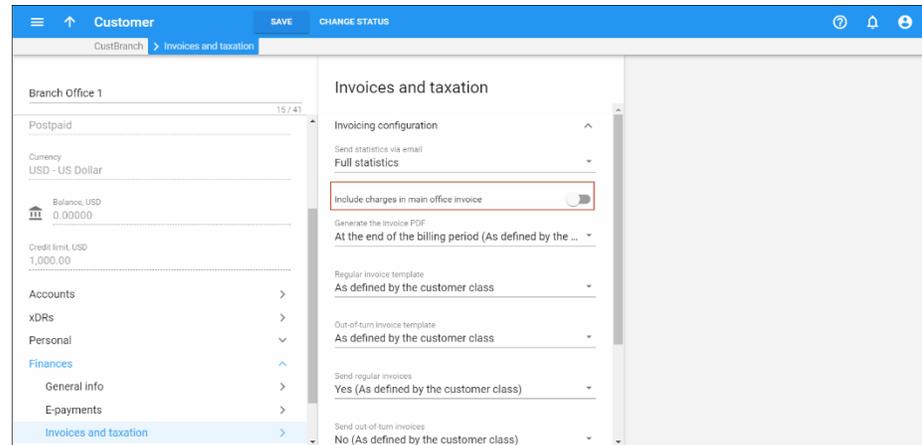
At the end of the billing period, the system generates a consolidated invoice for the Main Office (HQ) customer. This invoice includes the total charged amount for all of the company offices. The Main Office (HQ) customer also receives statistics reports generated for each branch office.

NOTE: To generate statistics reports for Branch Office (site) customers and send them to the Main Office (HQ) customer, the administrator must perform the following:

- Enable the **Send Stats If No Invoice** option on the Configuration server.
- Enable statistics generation for all Branch Office (site) customers in the group.
- Set up the Main Office (HQ) customer's email as the main contact for all Branch Office (site) customers in the group.

The system generates consolidated invoices in accordance with the time zone defined for the Main Office (HQ) customer. If Branch Office (site) customers are located in different time zones, you may need to postpone invoice generation until their billing periods are closed. In that case, use the customer class option **The billing period is closed**.

To generate an individual invoice for a branch office, the administrator disables the **Include charges into main office invoice** option for the corresponding Branch Office (site) customer.



Important notes:

- The administrator can configure an individual credit limit / spending plan for each linked customer in the group.
- If the Main Office (HQ) customer is blocked, suspended or has exceeded its individual credit limit, Branch Office (site) customers may continue using the services unless their individual credit limits / spending plans have been reached. At the end of the billing period a consolidated invoice is generated.
- If the Main Office (HQ) customer has been exported, or provisionally or permanently terminated, Branch Office (site) customers may continue using the services unless their individual credit limits / spending plans have been reached. At the end of the billing period, a consolidated invoice is not generated and the Branch Office (site) customers' charges are not moved to the Main Office (HQ) customer's balance.

This allows service providers to consolidate all of their communication charges on a single bill for their customers that have distributed infrastructure, and thereby provides greater convenience.

Customer termination

You may terminate a customer, including all his accounts. If, for some reason, you do not want a customer to remain in your PortaBilling® environment any longer, click the **Change Status** button on the **Customer Info** toolbar and then select the **Permanently Terminated** check box. This option allows you to stop all the customer's activities, and later to remove him and all his accounts from the environment. When terminated, the customer is no longer available for any operations. The only way to trace this customer is by using **Advanced Search** with the "Permanently Terminated" status filter.

Provisional customer termination

In addition to permanent termination you can use the *provisional termination* functionality.

Once a customer is provisionally terminated, all their services are closed (that is, no services can be used). But there is still an option to reactivate services that were disconnected if the customer should change their mind later on. The administrator is easily able to restore this customer's services so that they can be used exactly as before.

If a customer ultimately decides to discontinue services, the administrator can permanently terminate this customer in the system. In case of permanent termination, all customers' services are closed and cannot be restored.

Note: please keep in mind that if you are going to provisionally terminate a customer, this customer won't be charged for any DID numbers assigned to them. At the same time, the DID provider will still charge you a fee (e.g. \$5/month) because these DID numbers remain allocated to your network until the customer is permanently terminated.

Customer classes

Customer classes allow you to define which type of service a customer is using and a set of parameters which will be shared among a certain category of customers. For instance, suppose your invoice term for retail customers is "Net 21 days," while for business customers it is "Net 30 days." If your operators enter these values manually for each customer, there will inevitably be mistakes. Instead, you can define two separate customer classes, one named "Retail" and the other "Business," and define these parameters within them. After that, your operators need only assign a specific class to a given customer in order for the customer to automatically inherit all the class properties (grace period, invoice template, taxation, notification list and so on).

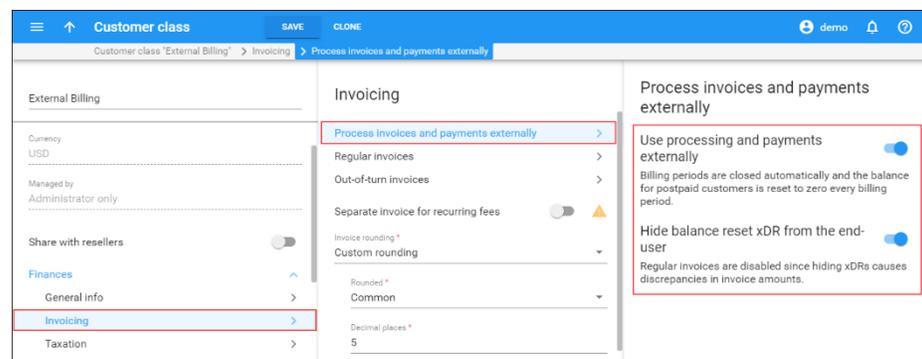
Balance control for customers billed externally

Service providers can maintain an actual balance for **postpaid** customers billed via external systems (i.e. when invoices are issued and processed via an external system).

This allows service providers to protect customers from fraud and unauthorized expenses by utilizing the credit limit tool. This tool restricts the amount of money that a customer can spend on services per billing period and suspends customer services when the credit limit is reached.

The administrator configures the following options for a customer class:

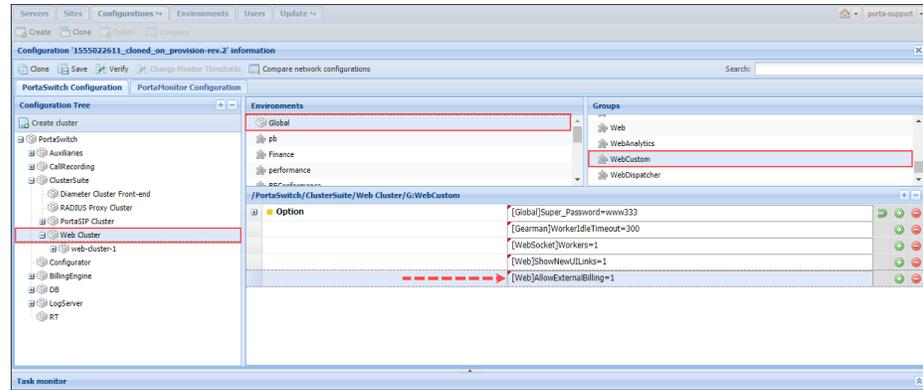
- **Use processing and payments externally** – This marks customers assigned with this customer class as being billed via an external system. Billing periods are closed automatically and the balance for postpaid customers is reset to zero by applying the manual credit transaction at the end of the billing period.
- **Hide balance reset xDR from end-user** – This defines whether or not to hide the reset balance transaction on the customer self-care interface. If this option is enabled, the regular invoices are disabled since hiding xDRs causes discrepancies in invoice amounts.



The administrator assigns this customer class to customers who are billed via an external system. At the end of the billing period, their balance is automatically reset to zero.

NOTE: Disabling the **Use processing and payments externally** option can cause discrepancies between balances and invoices. This could result in manual balance and invoice corrections needing to be made for each affected customer.

To enable this functionality, the administrator adds the `[Web]AllowExternalBilling=1` option for the **WebCustom** group on the Configuration server.



This allows service providers to use the credit limit tool for those of their customers who are billed via external systems, thereby protecting them from fraud and / or unauthorized expenses.

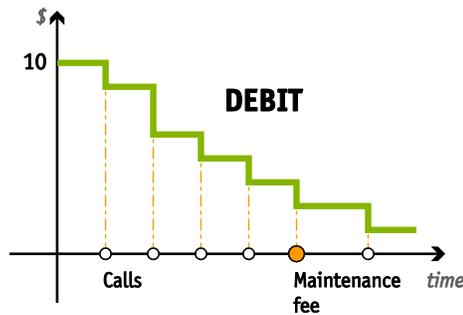
Account and account types

An account identifies the end user who is using the service. For example, a prepaid calling card is an account identified by a PIN, while a customer’s VoIP gateway could be an account identified by its IP address, and a SIP phone is identified by its phone number. There are two main types of accounts (debit and credit) plus an auxiliary one (voucher).

Debit

Typically used for: prepaid calling cards.

- The balance shows the available funds on the account. Initially, the balance is equal to the “opening balance” (typically the prepaid amount), and decreases with every call made.



- The account is unusable when the balance hits zero.

Credit

Typically used for: postpaid services.

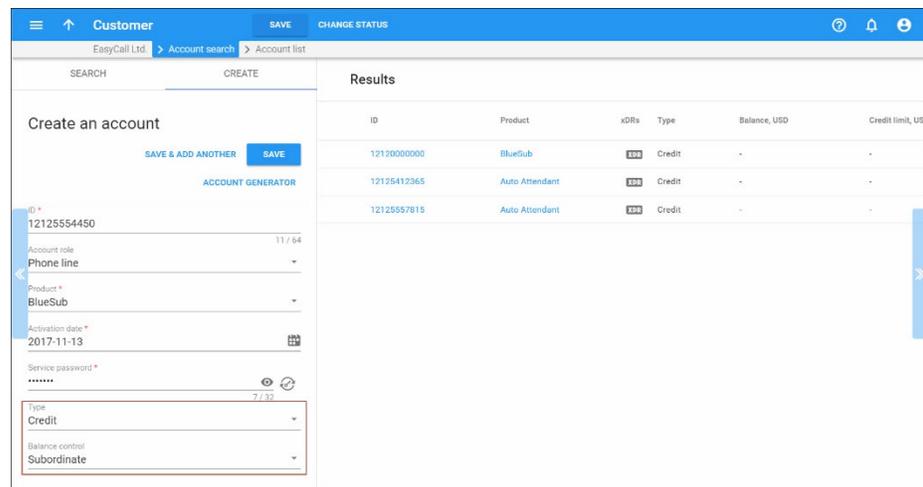
- Unlimited number of simultaneous sessions.

Subordinate credit accounts

In most cases, a customer makes a unified payment for all accounts and controls their own credit limit. Therefore, displaying the balance of individual accounts on the web interface is not necessary.

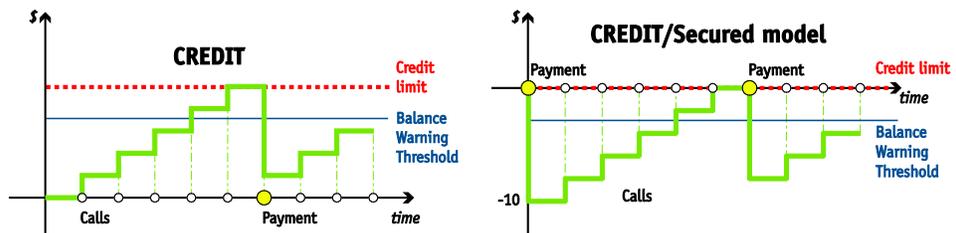
The administrator operates with the *customer's* balance only.

Consider the following example: a company has several phone lines (accounts). The users make calls, and therefore, the balance for the whole company increases. At the end of the billing period the company receives a consolidated invoice showing activities from all the accounts and sends a single payment which covers and is then applied to all accounts. Therefore, the company's balance decreases with one action.



Credit accounts with individual credit limit

- The balance reflects how much the account owner owes you. It starts with zero and goes up every time a service is used (e.g. a call is made, or a message is sent), and decreases when a payment is received.



- The account is unusable when the balance reaches the “credit limit”.
- The balance can also be negative – this means you owe money to the account owner (e.g. the account owner has made a deposit).

By setting the credit limit to zero you can provide “secured” services (as shown on the diagram above), in which case an account owner can only use a service when there is a deposit remaining in his account.

Voucher

Typically used for: refill debit or credit accounts for a customer using the same account over a long time, e.g. ANI-based authentication.

A voucher can be used to “recharge” (increase the available amount of) a credit or debit account. In this case, the balance of the voucher account is transferred to the main account. An account of this type cannot be used for any services such as making calls. Recharging may be done on the account self-care pages or via a special IVR script.

NOTE: End users can top up both a debit account and a credit account with an individual credit limit with a voucher that has a different product assigned. To top up an account with a sub-realm, a voucher must have the same sub-realm as the account being recharged.

Vouchers with encrypted PINs

Vouchers with encrypted PINs are composites. That is, they consist of visible voucher IDs and auto-generated passwords that are stored in encrypted form in the database.

Making balance top-ups using vouchers with encrypted PINs is done the same way as when using ordinary vouchers. When a voucher PIN is entered (either via the web self-care interface or the IVR application), the system decrypts the password, validates the voucher PIN and modifies the balance.



Note that previously generated vouchers cannot be used because their IDs do not contain encrypted passwords. To enable their usage, set *Yes* for the **AllowVouchersWithoutPassword** option on the Configuration server.

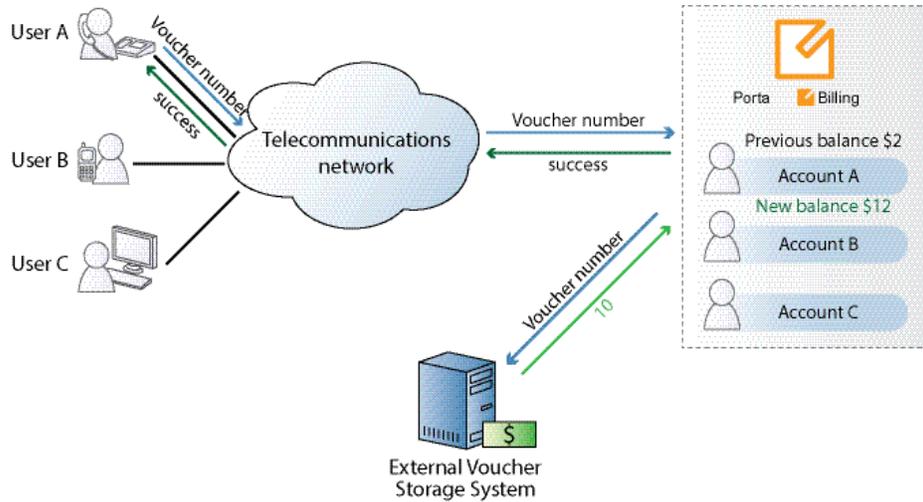
This enhancement enables service providers to introduce full voucher management within PortaBilling®, which additionally ensures business data security and complies with general system security demands.

Account top-up using external vouchers

For the most part, ITSPs use PortaSwitch® as an all-in-one solution in which all services (service provisioning, account management, invoicing, etc.) are provided solely by PortaSwitch®. However, ITSPs that partner with suppliers of top-up vouchers may prefer to retain the business processes they have established with those suppliers, in which case the

suppliers need to be integrated so they are recognized as part of the ITSP network.

PortaBilling® supports integration with external top-up voucher storage systems and retains all voucher management processes already established with those systems. This allows ITSPs to provide customers both with vouchers generated internally (i.e., by PortaBilling®) and vouchers generated externally (i.e., by their voucher suppliers).



When an end user requests account top-up by voucher using the IVR application or the account self-care interface, the voucher is first checked against the PortaBilling® accounts.

If the voucher is not found among the PortaBilling® accounts, it is considered external and the voucher number is sent to the external voucher storage system. There the voucher is checked for validity and the corresponding response is returned to PortaBilling®. The account balance is topped up based on the voucher face value and the voucher is marked as depleted on the external voucher storage system.



NOTE: The use of external vouchers is enabled on the Configuration server web interface.



Note that for account top-up with external vouchers the currency of the voucher is assumed to be the same as that of the customer.

Thus, by integrating external voucher storage systems with PortaBilling®, ITSPs are able to extend their partner relationships. Moreover, if their partners have large sales networks, integration may allow the ITSPs to leverage these networks to boost their own sales.

Account batches

Accounts can be grouped into batches. Each batch has its own descriptive name, giving you two ways to identify an account:

- By **account ID** (this is usually some type of ID the customer uses to access the service, e.g. PIN number for calling cards or phone number for residential VoIP services);
- By a combination of **batch name** and the account's **control number** (a sequential number within the batch).

When creating new accounts, you can either create them in a new batch or add them to an existing batch. In addition to better account tracking, this also permits **easy modification of a large number of accounts**.

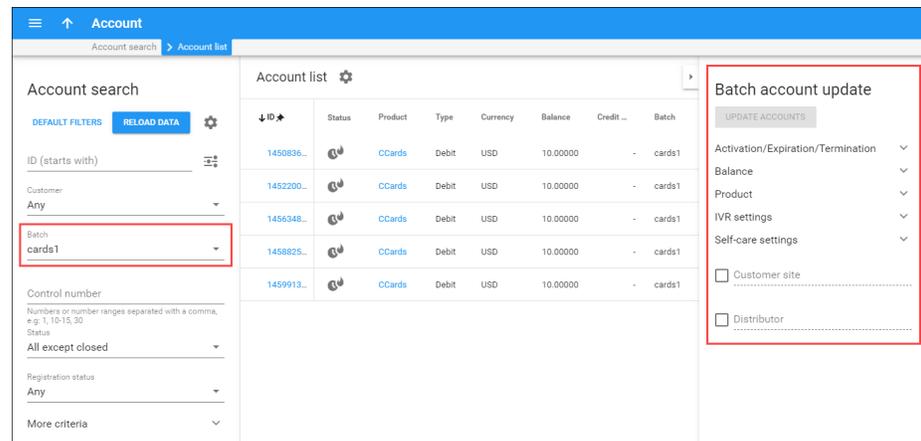
Note that a batch always belongs to a customer, so you cannot mix different customers' accounts within the same batch, and you cannot have two batches of the same name under two different customers.

You can also create new accounts without assigning them to any batch (this is usually done if you expect to have just a few accounts under a customer). You can always assign a batch to these accounts later.

Batch account update

The **Batch account update** function allows you to apply new settings to a group of accounts at a time. For example, you can assign another product to all the accounts within a batch at once.

To update multiple accounts, filter the batch of accounts you want to modify and specify the settings to update on the **Batch account update** panel. You can activate or block accounts, change their balance, assign products or change other settings for either the whole batch or for specific accounts.



The screenshot displays the 'Account list' page in the Porta Billing system. On the left, the 'Account search' sidebar includes filters for ID, Customer, Batch, Control number, Registration status, and More criteria. The 'Batch' filter is set to 'cards1'. The main area shows a table of accounts with columns for ID, Status, Product, Type, Currency, Balance, Credit, and Batch. On the right, the 'Batch account update' panel is visible, featuring an 'UPDATE ACCOUNTS' button and a list of settings to be updated, including Activation/Expiration/Termination, Balance, Product, IVR settings, Self-care settings, Customer site, and Distributor.

To update specific accounts in a batch, you can apply additional filtering options such as control number, registration status, etc. Note that filtering by custom fields is not supported for batch account update.

You can see the list of applied filtering criteria and settings to be changed in the confirmation dialog window once you click the Update accounts button.

Account life cycle

An account may be created in the inactive state. This enables you to fully provision the entire service configuration for this account, yet the service can only be used once the account is activated.

There are two parameters that define the life cycle of an account:

- Every account has an **activation date**. This defines the date of the account first usage (e.g. you print Christmas promotion cards in November, but set the activation date for December 23rd so that no one can start using them earlier). Note that an account is activated on a specific day in the customer's billing time zone – the time zone in which the customer's billing periods close (e.g. if you are in Sydney and your customer is in New York, the Christmas promotion cards you have created will be activated on December 23rd New York time).
- **Expiration date** defines the date after which the account can no longer be used. There are three options for setting it up:
 - **Specific Date** – For instance, the administrator sets March 20th as the day of expiration, so the end user may use the services till 23:59 on March 19th. After that, the account is expired.
 - **Account availability after the first use** – The customer assigns how many days the services will be available for after the first use. For instance, the period for using the calling card service was defined as 30 days. An end user

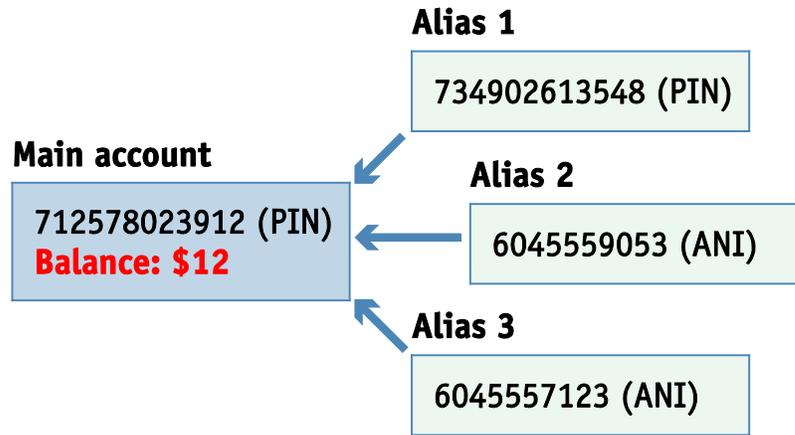
- makes the first call on the 1st of March at 11:00 and uses the calling card daily. Therefore, the 30th of March 23:59 is the last time when the end user can use this calling card.
- **Account expiration after last use** – The administrator assigns after how many days from the last use the services will be available for. For instance, this period equals 30 days. The end user makes his first call on the 21st of February. He also makes calls on the 23^d, 25th of February and on the 1st of March at 11:00. Therefore, the 30th of March 23:59 is the last time when the end user can use this calling card.

When configuring an account expiration date, either select one or combine several expiration options. They take equal priority: an account expires according to which event occurs first. Additionally, the expiration date is optional, so you may create accounts that have no expiration date.

Another important point regarding expiration options is that a voucher recharge operation extends the life cycle of an account. Thus, if an account is adjusted to be available for 60 days after first use, and the account's owner recharges the account 57 days after the date of its first use, the service is extended for another 60 days. The account recharge works analogically for account expiration after last use.

Account aliases

Sometimes it is desirable to provide multiple means of authentication for the same account; for instance, several phone numbers may be registered for PINless dialing to a single prepaid card. In this case, there will be a single account containing the actual balance and other billing information, and multiple account aliases associated with it. An alias serves as a link or pointer to the main account. When a customer uses an alias ID for authentication, PortaBilling® retrieves information about the main account and uses its balance, product and other parameters for all further operations.



As mentioned above, an alias gives you the ability to use the service via an additional identification, while still drawing on funds from the main account. In some cases you may want to restrict this ability; one typical example is DID forwarding in the case of the SIP trunking service. Here independent registrations for each additional alias created under the main account will be prevented (this is done by clearing the corresponding check box in the alias configuration or disabling the alias authentication option at the customer level). In this case, incoming calls to the alias will be routed to the main account. Outgoing calls using the alias cannot be made.

Customer and account statuses

Customer statuses

Please note that when no status is displayed, the customer is active and their services are available. The status displayed depends upon priority. In the list below, the status is indicated in the Priority column. Number 1 corresponds to the highest priority.

A customer may have the following status:

Priority	Status	Description
1	 Closed	The customer’s account has been closed and is only in the database for informational / regulatory purposes. No further operations are possible with this entity.
2	 Blocked	The customer has been blocked by the administrator; no call services are provided until the administrator removes the block. The customer has no access to self-care pages. Blocked customers are

		not subject to maintenance charges but subscription fees still apply.
3	 Suspended	Services for this customer have been suspended because of an overdue invoice. Once payment is received (either an online payment or periodical payment, or the administrator enters a cash / check payment manually) in an amount that covers the overdue invoice, the suspension is automatically lifted.
4	 Provisionally terminated	Once a customer is provisionally terminated, all services are terminated. There still exists the option to reactivate the services that were disconnected if the customer changes their mind. The default period for reactivation is 30 days, however this value can be adjusted within the customer class.
5	 Credit exceeded	This indicates that the postpaid customer's balance is above his credit limit, so he will not be allowed to make any outgoing calls unless they are free calls (e.g. toll-free numbers).
6	 No available funds	This status is only for prepaid customers. When a customer's funds reach a zero value, services are no longer available. Once a payment is made and available funds increase, all services are once again available.
7	 Suspension lifted	This indicates that the customer's suspension has been lifted (the 'as of' date is specified on the Change status button).
8	 Payment frozen	This indicates that the customer's auto-payment has been suspended due to repeated errors (e.g. credit card cancellation).
9	 Spending limit reached	This indicates that the customer's daily spending limit has been reached. The spending limit is renewed at 00:00:00 every day, so once the limit is renewed, the Spending Limit Reached status is lifted from the customer.
10	 Exported	This status indicates that the customer has been ported to a new PortaOne® installation. When an account has this exported status, it is not operational

		although its data is preserved for backup.
--	--	--

Customer statuses and service availability

In the tables below you can see whether services are allowed or denied depending on a customer status, a type of service and overdraft protection settings.

- 1) The **Overdraft Protection** setting for the accounts' product is the following:
 - The **Remaining balance requirement** option is set to **No Restriction** (the default behavior).

Customer Status	Type of service	
	Toll-free	Chargeable
Credit exceeded / No available funds	allowed NOTE: The toll-free services are denied for accounts under reseller customers if the reseller's credit limit is exceeded.	denied NOTE: The chargeable services are allowed for debit accounts.
Spending limit reached	allowed	denied
Blocked	denied	denied
Suspended	Service availability depends on the Service Allowed to Suspended Users option. When set to None , no services are provided to end users. When set to Zero Charged Only , only toll-free services are available.	
Suspension lifted	denied	denied
Payment frozen	allowed	allowed
Provisionally terminated	denied	denied
Closed	denied	denied
Exported	denied	denied

- 2) The **Overdraft Protection** setting for the accounts' product is the following:
 - The **Remaining balance requirement** option is set to **Positive amount available**.

Customer Status	Type of service	
	Toll-free	Chargeable

Credit exceeded / No available funds	denied	denied
Spending limit reached	denied	denied
Blocked	denied	denied
Suspended	Service availability depends on the Service Allowed to Suspended Users option. When set to None , no services are provided to end users. When set to Zero Charged Only , only toll-free services are available.	
Suspension lifted	denied	denied
Payment frozen	allowed	allowed
Provisionally terminated	denied	denied
Closed	denied	denied
Exported	denied	denied

Reseller statuses affect service availability in the same way, assuming that accounts do not inherit additional statuses from their customer (i.e. the customer is active and can use the services).

Account statuses

Please note that only one status is displayed in an account although the account may inherit additional statuses from the customer. The status displayed depends upon priority. In the list below, the status is indicated in the Priority column. Number 1 corresponds to the highest priority.

An account may have one of the following statuses:

Priority	Status	Description
1	 Expired	An expired account is no longer active and cannot be used.
2	 Screened	This indicates that suspicious activity has been detected. If an account makes a call from what a service provider defines as a high-risk location, the status is changed to Screened. The account is then connected to a screening IVR where additional credentials must be provided to prove that this is, indeed, a legitimate user. On-net calls, however, are still available.

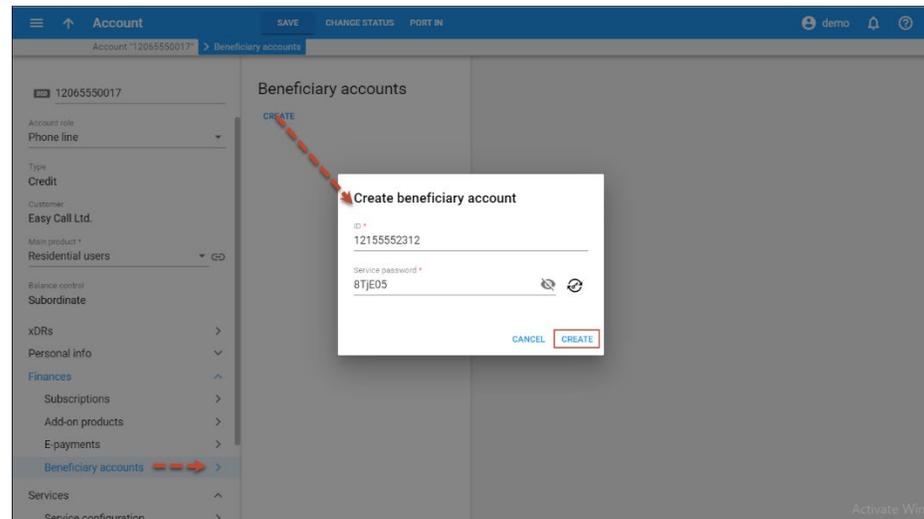
3	 Quarantined	<p>This means that after being screened, this account was unable to supply valid credentials and is still generating many call attempts. In order to reduce the load on the network, all off-net call attempts are automatically blocked. However, on-net calls are still available.</p>
4	 Closed	<p>The account has been closed and is only in the database for informational / regulatory purposes. No further activities are available for this entity.</p>
5	 Inactive	<p>All debit accounts with Prepaid cards, Vouchers, or with assigned SIM cards that are created via the Account Generator begin with this inactive status to prevent misuse of services. The accounts can only be used to access the service once they are manually activated.</p>
6	 Suspended	<p>This status is inherited from the customer. Services for this account have been suspended due to customer non-payment. Once payment of an amount that covers the overdue invoice is received, the suspension is automatically lifted.</p>
7	 Customer provisionally terminated	<p>This status is inherited from the customer. Once a customer is provisionally terminated, services are unavailable for their accounts. The option to reactivate disconnected services so that they are available as before still exists.</p>
8	 Blocked	<p>The account has been manually blocked by an administrator; no services are provided until an administrator removes the block. The account has no access to its self-care pages. Blocked accounts are not subject to maintenance charges though subscription fees still apply.</p>
9	 Customer blocked	<p>This status is inherited from the customer. The customer has been manually blocked by an administrator; no services are available until an administrator removes the block. The account has no access to its self-care pages. Blocked accounts are not subject to maintenance charges though</p>

		subscription fees still apply.
10	 Not yet active	This status indicates that the account's activation date has not yet occurred. The account will be automatically activated on a future predefined date. While the account is inactive, activities are unavailable.
11	 Credit exceeded	This indicates that the account's balance has gone above the individual credit limit and services are no longer available. Once payment is made, all services are once again available.
12	 Customer credit exceeded	This indicates that the customer's balance is above the credit limit so services are no longer available. Once payment is made, all services are once again available.
13	 Overdraft	Overdraft status is assigned to debit accounts only. Once an account's balance is overdrawn, services are no longer available. Once payment is made and available funds increase, all services are again available.
14	 Customer has no available funds	This status is inherited from the prepaid customer. Only subordinate credit accounts may have this status. When customer's funds reach a zero value, services are no longer available. Once a payment is made and available funds increase, all services are once again available.
15	 Zero balance	Only debit accounts may have this status. When funds reach a zero value, services are no longer available. Once a payment is made and available funds increase, all services are again available.
16	 Suspension lifted	This status comes only from a customer. This status indicates that the customer's account suspension has been lifted.
17	 Exported	This status indicates that the account has been ported to a new PortaOne® installation. When an account has this exported status, it is not operational although its data is preserved for backup.

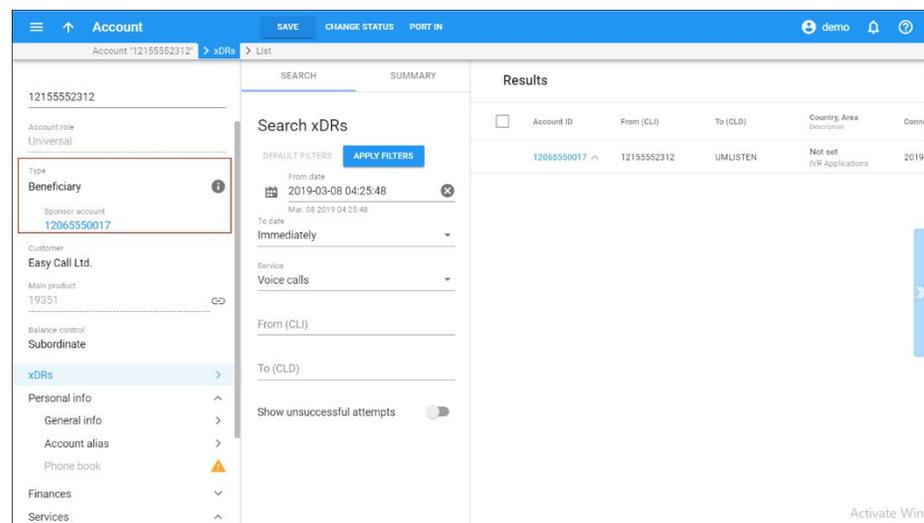
When no status is displayed, the account is active and services are available.

Beneficiary accounts

You can separate service and billing configuration for particular accounts. Accounts with only a service configuration are called beneficiary accounts. Beneficiary accounts are linked to regular accounts and inherit their billing configuration – balances, quotas and tariff rates. In other words, a regular account “sponsors” its beneficiaries.



When a beneficiary account uses the service, charges and quota consumption apply to the sponsor. Likewise, the xDRs are created for the sponsor although the beneficiary is the actual user of the service.



In contrast with using account aliases, beneficiary accounts can have an individual service configuration (e.g. their own SIM card, individual

extension or different service policy). You can only configure the service features for the services included in the sponsor’s main product.

A customer can access the self-care portal for their beneficiary accounts to configure the individual service features. For example, to configure the auto-attendant on a beneficiary account’s self-care portal.

Beneficiary account owners can also access their self-care portal where they can edit their personal info, IVR language or change their call forwarding settings at any time.



An administrator can add and manage beneficiary accounts via both the web interface and the API. Resellers can currently operate with beneficiary accounts via the API only.

Sponsor and beneficiary accounts’ activities

The table below compares the sponsor and beneficiary accounts in terms of their activity:

Activity	Sponsor accounts	Beneficiary accounts
Account configuration	Any account automatically becomes a sponsor as soon as a beneficiary account is created under it. To switch off sponsorship for the account, terminate all of its beneficiaries.	A beneficiary account can only be created from scratch under a regular account. You cannot convert a beneficiary account into a regular one or make it a sponsor.
Service configuration	Can have an individual service configuration that is different from the beneficiaries.	Can have an individual service configuration (e.g. enable voicemail or assign a different access policy). You can override the service features at the account level.
Product configuration change	Allowed. You can change / modify the main product (e.g. add a new service or change the tariff), assign / remove add-	Forbidden as the beneficiary account inherits its product configuration from the sponsor.

	ons.	
Volume discount plan usage	Shares its volume discount plans with all its beneficiaries. You can add volume discounts, quotas and service wallets.	Uses the volume discount plans assigned to the sponsor. When a volume discount plan is finished or the CLD is not included in the destination group set, charges apply to the sponsor.
Subscription assignment	Allowed via add-ons and also directly. All recurring charges apply to the sponsor.	Restricted. The sponsor account bears the recurring charges of its beneficiaries.
xDR history	Includes its own xDRs and xDRs created on behalf of all beneficiaries. These xDRs contain the beneficiary account ID.	Shows the sponsor's xDRs created on behalf of a beneficiary account.
Balance and payments	Sponsor accounts see own balance / available funds. Account owners can make credit card payments and top ups via vouchers.	Beneficiary accounts see their sponsors' balance. Account owners are forbidden to make payments and top ups.
Life cycle	Independent from beneficiaries. When a sponsor account expires, services for beneficiaries are blocked.	Independent from sponsor. The date of last usage is updated for both sponsor and beneficiary accounts. The account's expiration does not affect service availability for the sponsor.
Status change	Any status change affects service availability for both sponsor and beneficiary accounts.	The status of the beneficiary account does not affect its sponsor yet the beneficiary account is

	For example, when a sponsor is blocked, services are denied both for it and for its beneficiaries.	subject to the sponsor's status change.
Self-care access	Allowed.	Allowed.
API access under account realm	Allowed.	Allowed for the activities of a beneficiary account.

With beneficiary accounts you can offer shared buckets to your business customers and let them optimize their service consumption. For example, let your business customers share quotas among a certain set of their accounts (e.g. grant 500 minutes of international calls for all employees in the Sales department while the Support department receives 100 minutes). To meet the customer's request, your administrator creates a sponsor account for the Sales department, assigns a 500-minute quota to it and then creates the appropriate number of beneficiary accounts under it. Likewise, your administrator creates sponsor and beneficiary accounts for the Support department with the 100-minute shared quota assigned to the sponsor. As a result, customers optimize their service consumption and their customer experience improves.

Customer business models and account roles

It is essential for every service provider to have a thorough understanding of their customers, e.g. which services they use or would like to have, and what their needs are, etc. In order to simplify customer management in a convergent environment where different types of services are provided, PortaBilling® introduces two important concepts: business model and account role.

Business model

A business model defines which type of service a customer is using. The following business models are available:

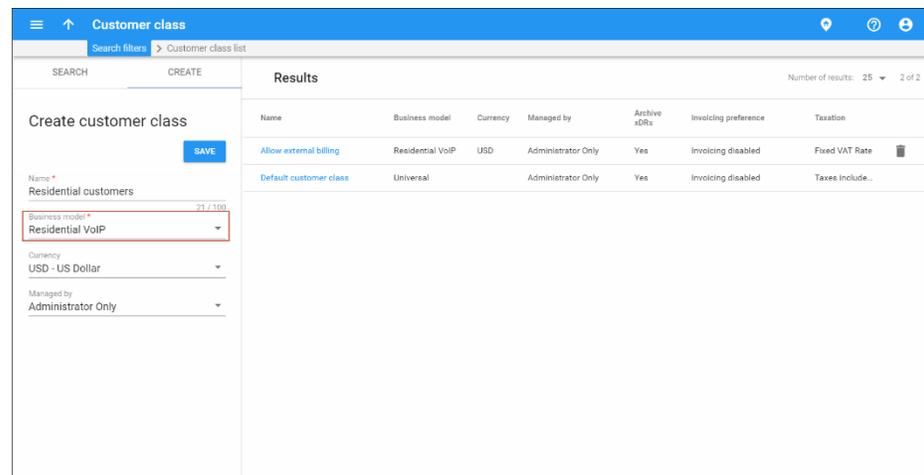
- Hosted IP PBX
- SIP trunking
- Residential VoIP
- PINless
- Internet access
- Prepaid cards holder
- Mobile
- Universal

The first six business models in the list are self-explanatory, while the other two must be explained in more detail.

The Mobile business model is customized for configuring mobile network subscribers – customers who are provided with a full range of mobile services.

The Universal business model is a generalized model that allows you to configure any type of service for your customers. In addition, the Universal business model is also used for backward compatibility – any previously created customer is automatically assigned the Universal business model after a software upgrade.

An administrator assigns a business model via a customer class.



Depending on the chosen business model, the administrator is offered a list of recommended account roles for it.

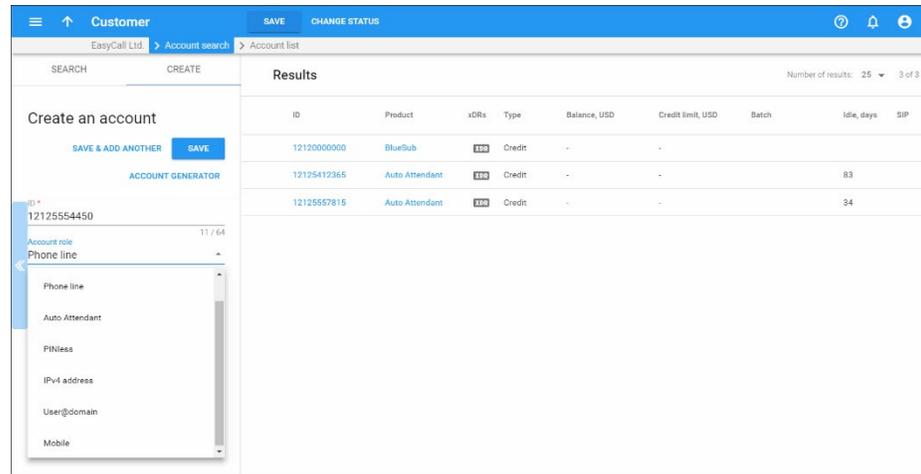
Account role

An account role defines what a specific account is designated for (e.g. whether the account represents a phone line or a top-up voucher) and executes account ID validation. Thus, for a phone line, one can only pick a valid phone number as an account ID.

The following table demonstrates which account roles are available for each business model:

Business model	Account roles
Universal	All available account roles.
Hosted IP PBX	Phone line, Auto attendant.
SIP trunking	Phone line, IPv4 address.

Residential VoIP	Phone line, Auto attendant, Voucher.
PINless	PINless.
Internet access	User@domain, Voucher.
Prepaid card holder	Prepaid card, PINless, Voucher.
Mobile	Mobile, Voucher.



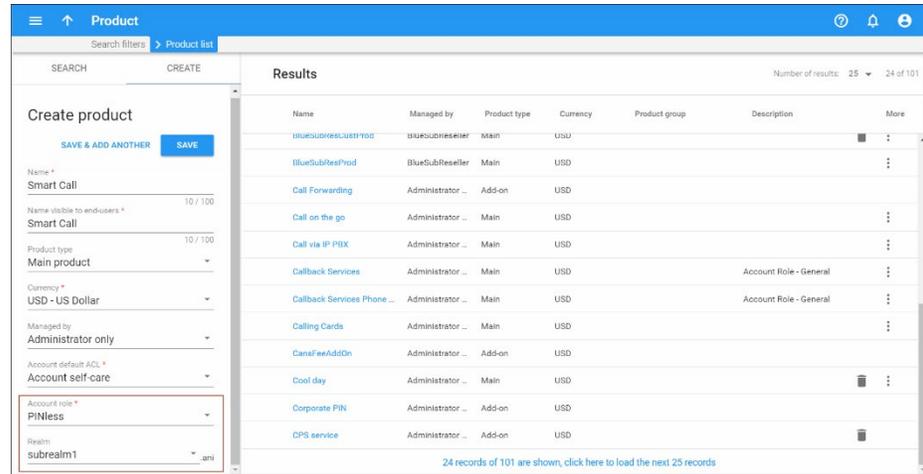
The administrator defines account roles for both accounts and products. This introduces the following constraint – a product can only be assigned to an account having the same role.

This makes it possible to reduce the occurrence of errors during product assignment. For example, an account that is supposed to be for a phone line can only be assigned to a product that is designated for that; i.e. one cannot mistakenly pick a business SIP trunk product. The administrator can change an account role if required.

Along with an account role, an administrator can define an account realm.

Account realm

An account realm allows administrators to impose a scope of uniqueness for an account ID. An account realm is configured for a product.



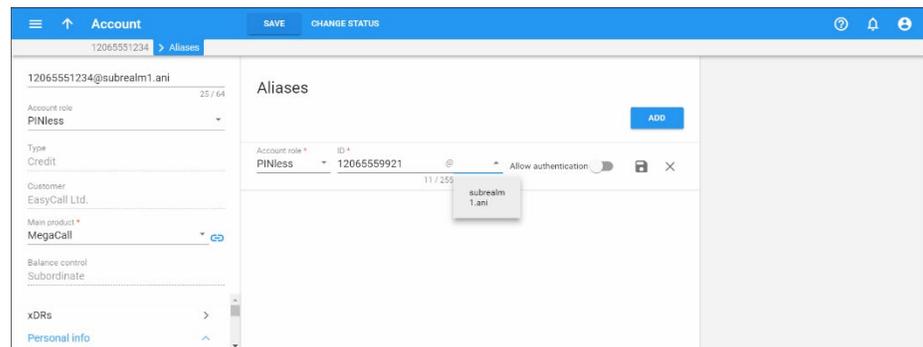
For example, user John Doe wants to subscribe to the SmartCall product for PINless dialing from their mobile phone 12065551234. To provision John Doe’s mobile phone in PortaBilling®, the administrator configures the SmartCall product using an **ani** realm. The administrator adds an account and chooses the SmartCall product for it. Then the account is saved in the following format: 12065551234@ani.

John Doe finds another PINless dialing product attractive, MegaCall, and wants to subscribe to it too. The administrator needs to add an account that must not be confused with the account for the SmartCall product. The account realm perfectly serves this purpose. The administrator configures the MegaCall product with **subrealm1.ani** realm. The administrator adds an account and chooses the MegaCall product for it. The account is saved in the following format: 12065551234@subrealm1.ani.

This allows proper provisioning for both accounts in PortaBilling®.

Alias role and realm

When an administrator adds an alias to an account, a role and realm must be selected for it. Alias role and realm selection is limited by the main and add-on products assigned to the account.



NOTE: The **Voucher** role is not supported for account aliases.

The introduction of the business model and account role concept ensures that:

- The correct account ID format is preserved during account creation; and
- The chances of assigning an incorrect product to an account are significantly reduced.

This makes the account creation procedure more user-friendly and reduces the occurrence of human error.

Distributors

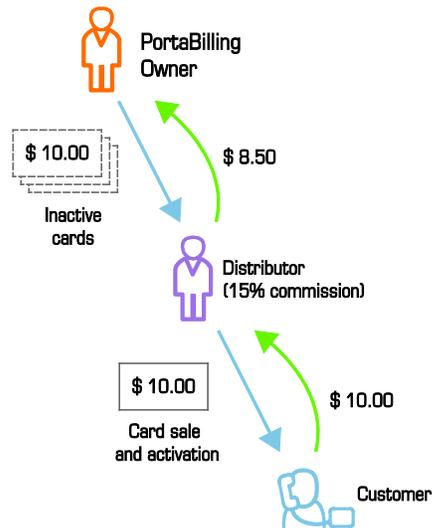
The distributor model is designed to expand sales activities by engaging additional agents and enlarging the point-of-sale network, without any significant costs or risks.

Customers can purchase new products, settle their invoices or refill their prepaid accounts by paying cash to a distributor (for example, the owner of the grocery shop opposite your house could be a distributor). After the distributor collects the money, he delivers it to the ITSP, minus his commission. PortaBilling® automatically calculates the commission and checks the distributor's balance to keep track of how much he owes the ITSP, and also to verify the distributor's credit limit. This helps to avoid a situation where the distributor would activate too many accounts without first submitting payment for accounts already activated, and thus limits the ITSP's risk of loss in case the distributor goes out of business. A distributor need not have any special technical knowledge or skills. They either deliver to the end user a tangible product (e.g., SIM card, calling card) supplied to them beforehand in bulk, or, in the case of an intangible product (e.g. registering a customer's cell phone number so they can make PINless dial calls), they can create accounts/customers using a quick form. All account management operations are done via the web interface.

While a distributor might only sell a few new accounts each day, creating accounts in such small amounts at the distributor's request requires too much administrative overhead. Also, when prepaid cards or top-up vouchers are being printed in the print shop, it makes sense to do this for many thousands of cards at once.

Therefore, the account generator can be used to produce a large batch of accounts so that all the cards can be printed at once. Then later a distributor can be assigned to some of these accounts (i.e. when this distributor receives a portion of cards to sell).

Nonetheless, when some of a batch of cards is delivered to a distributor their total face value could be quite large. So it is risky to allow all of the cards to be active (e.g. if the cards are stolen, they could all be used). Thus cards (accounts) are typically generated and supplied to the distributor in an **inactive** state, and only when the distributor activates the card (account) it can be used to access the service. Likewise, the distributor is only charged once activation is complete.



The following are associated with a distributor:

- **Default Sales Commission** (applied when an account is activated or a distributor is assigned to an active account);
- **Default Payment Commission** (applied when a payment is entered by a distributor).

When a customer or account is activated or created using a quick form under a distributor, the latter's balance increases by the account's balance, minus his commission. For example, if the distributor's commission percentage is 15% (the default sales commission) and he sells a new account with a balance of \$10, the distributor is charged \$8.50 (this is the amount he owes to the ITSP). So the distributor receives \$10 from the customer and gives back \$8.50 to the ITSP, thus making a profit of \$1.50. If the distributor applies payment of \$10 towards an account, and his payment commission percentage is 10% (the default payment commission), he is charged \$9 and receives a \$1 profit.

Thus the distributor is charged when:

- he applies payment towards a customer or an account;
- he is assigned to an active account;
- an account to which he is assigned has been activated, or a new account is created in the active state.

To prevent operator error when attempting to generate accounts being assigned to a specific distributor, the account generator will only allow the creation of inactive debit accounts/vouchers, and these accounts will have to be activated later.

Typically used for: Point-of-sale agents, where an agent delivers the product to end users, collects payments, and then delivers the money (minus his commission) to the service provider.

Multi-level distributors

Two distributor models are supported: the “simple” and the multi-level. The multi-level distributor model allows ITSPs to further expand their sales activities and the size of their distribution network to increase their customer base and total revenue.

In the multi-level model, distributors are able to create their own subdistributors with whom they share the part of their commissions. The amount of a subdistributor’s commission cannot be higher than the one the service provider defines for the distributor.

NOTE: Distributors can only create subdistributors in the same currency that they operate in. For example, a distributor who operates in USD can only create a subdistributor in USD.

While distributors operate with end users and with subdistributors, subdistributors operate with end users only.

Before a transaction takes place between an end user and a subdistributor, the system checks whether the parties’ balances are sufficient to cover the transaction. This ensures that credits and balances are controlled efficiently within this segment of the distribution chain.

The following example of commission-sharing by a distributor and the distributor’s subdistributor illustrates how this functionality works.

An administrator creates Distributor X, assigning the default distributor payment commission of 20%. Distributor X logs into the distributor self-care interface to create Subdistributor Y, assigning the default subdistributor payment commission of 10%.

Subdistributor Y logs onto the subdistributor self-care interface to create retail customers. The subdistributor then receives a \$10 cash payment from a retail customer and registers a manual payment of \$10 on the customer’s behalf to the service provider of record (either the distributor or the ITSP). As a result, the balances of the parties involved change as follows:

- The customer’s balance is topped-up for a \$10 to be spent on services.

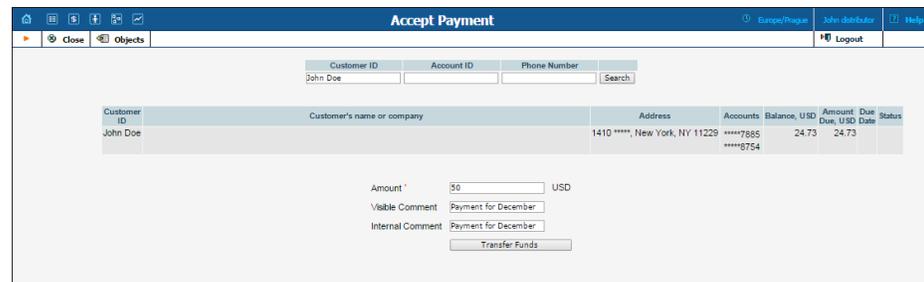
- Subdistributor Y owes the Distributor X \$9 (customer payment amount less the subdistributor commission: $\$10 - (\$10 * 10\%) = \$9$).
- Distributor X owes the service provider \$8 (customer payment amount less the distributor commission: $\$10 - (\$10 * 20\%) = \$8$).

In this way, multi-level distribution allows ITSPs to build a branched distribution network and thereby increase total revenue.

Payment recording by a distributor

Apart from creating customers and / or activating accounts, distributors' functions also lie in recording payments made by end users to top up their balance. When end users do not have the means to make payments online (e.g. they have no credit card), they can make cash payments via distributors. Distributors record such payments to PortaBilling®.

Distributors can record payments for any “direct” customer (i.e. a customer who does not belong to any reseller). However, to prevent distributors from acquiring customer information and misusing it, they can only search for customers and accounts if they enter either a full customer / account ID or a full phone number. Then they will only see limited information about a customer or an account.



When a payment has been accepted and recorded to PortaBilling®, a distributor is charged the payment sum minus his commission (e.g. the distributor applies a \$10 payment towards an account, and with his 10% commission, the distributor is charged \$9 and earns \$1 in profit).

This rule applies both to the customers this distributor is assigned to and to any other direct customers of an ITSP.

Distributor xDR history

Distributors can view only their own xDR history, which consists of the transactions they made. Thereby ITSPs prevent distributors from acquiring customer and account information and misusing it.

Distributors can reverse only payment transactions, e.g. if they have entered the wrong sum when recording the customer’s payment. The xDRs produced for account activation are not reversible.

Distributor commission report

For each transaction performed by a distributor, an xDR is created and the commission is calculated. However, showing the distributor’s actual profit could be more advantageous.

Distributors have the ability to track the profits they gain from sales and payment collections and the relative commission percentage and commission amount are provided in the form of a report.

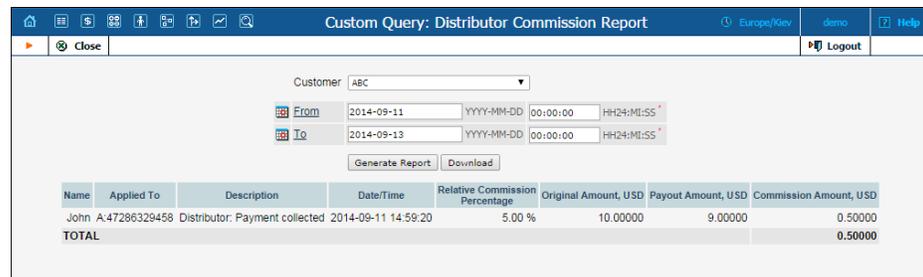
When enabled on the Configuration server, commission reports are accessible for distributors via the self-care interface.

Consider the following example:

The administrator creates the ABC distributor with a 20% commission. The ABC distributor creates the John subdistributor with a 10% commission.

The John subdistributor makes a manual payment of \$10 for the retail customer.

The John subdistributor launches the commission report where the commission percentage (10%), original amount of funds (\$10), payout amount of funds (\$9) and commission amount (\$1) are declared.



Name	Applied To	Description	Date/Time	Relative Commission Percentage	Original Amount, USD	Payout Amount, USD	Commission Amount, USD
John A:47286329458	Distributor: Payment collected		2014-09-11 14:59:20	5.00 %	10.00000	9.00000	0.50000
TOTAL							0.50000

NOTE: As the operation was performed by the John subdistributor, the commission percentage for the ABC distributor is the difference between ABC’s total commission (20%) and the commission assigned to John (10%).

The commission report for the ABC distributor reflects the operation performed by the John subdistributor with the ABC distributor’s commission applied: the commission percentage 10%, original amount of funds \$10, payout amount of funds \$8 (i.e. \$10 - 10% - 10%) and the \$1 commission amount.

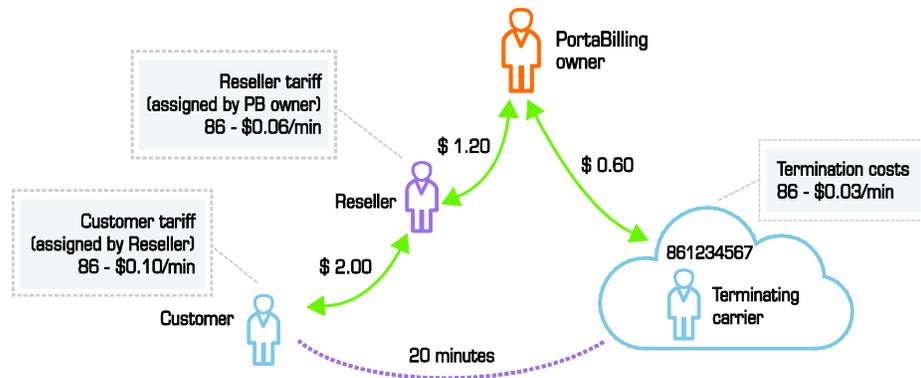
Thus, using commission report tools, distributors have a clear idea of the profits received from performing business operations.

The ITSPs are able to receive general information about all the distributors or only for a particular one.

Resellers

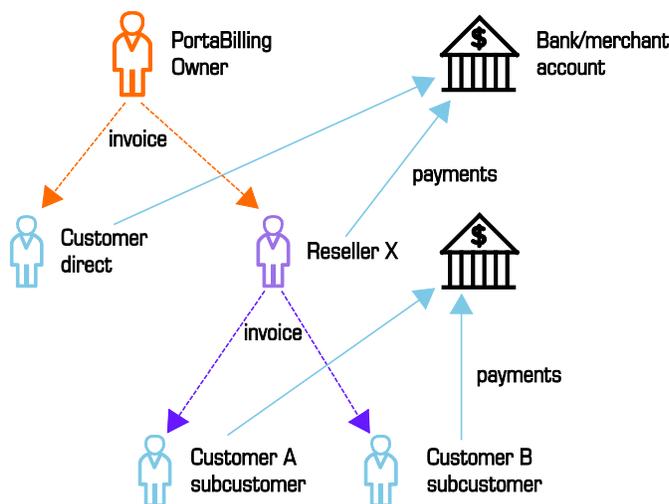
A reseller is a company that operates using your network and billing system, but sells the service under its own name. So while this company is presented as an independent operator to consumers, in reality it does not have any equipment of its own; the company registers its own retail customers (also called subcustomers) and individual accounts in PortaBilling®, and the billing of their end users is done there. This enables you to expand your operation by means of “white label” or “virtual operator” services.

When a call is made by an account (end user), the reseller is billed using his own tariff (wholesale rates). Thus the end user (account+retail customer to whom it belongs) and reseller are billed separately; they may even be in different currencies. So if, for instance, a partner is reselling your services, he will pay you \$0.06/min. for each call made to the US by his subscribers. At the same time, he can charge his customers any rate he wishes (e.g. \$0.10/min).



Resellers can perform most everyday operations such as managing their products, changing rates to their subscribers via the self-care interface, or browsing xDR records, so there is less workload for you as an ITSP. PortaBilling® will automatically generate invoices for a reseller’s customers with his name on them, and subcustomers can go to a specific domain to access their self-care interface, where they will see the reseller’s name and logo. This ensures that the end user has an image of the reseller as an independent operator, and is not aware of your role as the ITSP.

Reseller collects money from his subcustomers and is responsible for paying invoices, issued by the PortaBilling® owner. If reseller exceeds his credit limit – this automatically blocks all activities of his subcustomers.



Typically used for: OEM or “white-label” service providers, companies who would like to offer the service under their own brand name, while the actual service is provided on your VoIP network and all the billing and management of their customers is done in PortaBilling®.

NOTE: Billing via the reseller model implies an additional load on the billing engine, compared to a retail customer, and using only retail customers is about 50% faster than when a reseller is involved. If performance is an issue, avoid creating resellers when it is not necessary.

Multi-level resellers

To reach more customers in the market many ITSPs sell their services through a multi-level network of resellers. To implement this functionality in PortaBilling®, the “carrier – reseller – end user” model has been extended to a multi-level reseller structure where the entire chain of resellers is virtually unlimited.

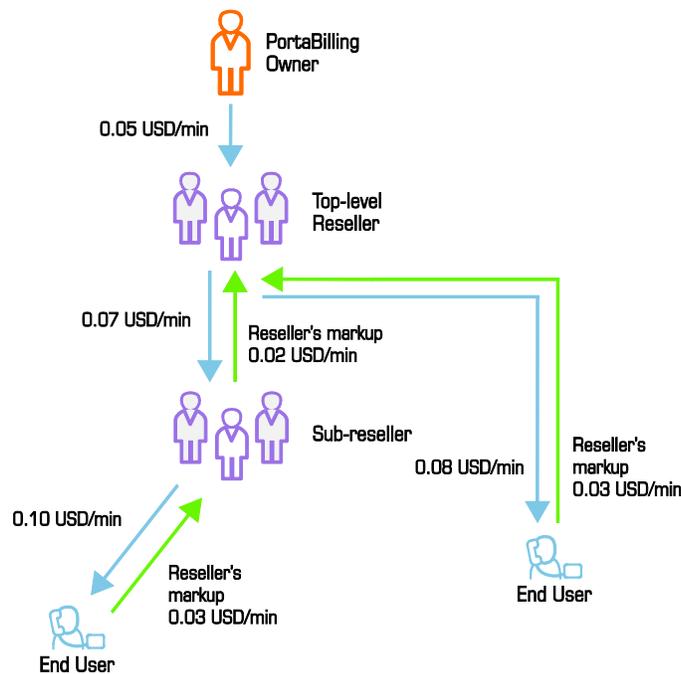
Let’s assume that an ITSP engages Reseller A to sell VoIP services. Reseller A operates independently from the ITSP and is free to set his own prices for services provided. Reseller A begins to sell services to his direct clients and some small business owners thus building his own network of subresellers. Each subreseller engaged by Reseller A can act as a “white label” reseller and resell a package of services (e.g. through their own network of reselling partners) under their own brand name.

In this scenario, it might appear that services purchased by customer John Doe from some reseller have in fact been provided through the n-level

network of resellers. Consider the example of the resellers' network shown in the diagram below.

The reseller's hierarchy is modeled as follows:

- Top-level reseller – This reseller stays at the top of the resellers' hierarchy. He is created and managed by an administrator and deals directly with the service provider. He can create and manage his own subresellers as well as provide services directly to end users.
- Subreseller – The subreseller is created and managed by a top-level reseller or another (higher-level) subreseller. Likewise the top-level reseller, the subreseller can create and manage their own (lower-level) subresellers and provide services to some end users.



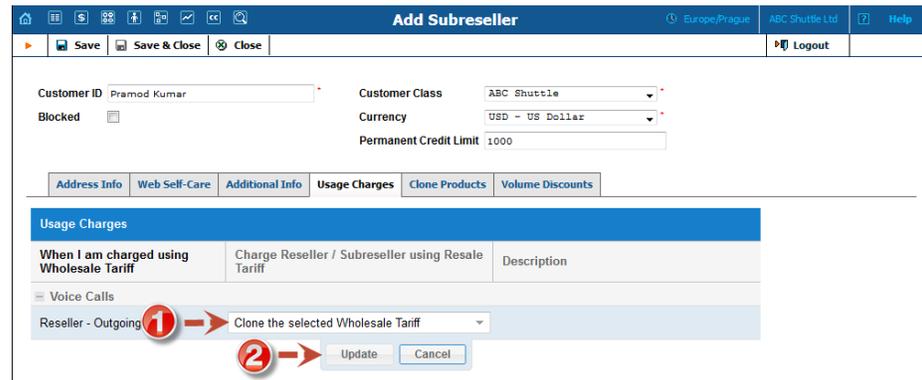
Managing tariffs

Each reseller operates with three types of tariffs:

- Wholesale tariffs – are tariffs according to which the reseller is billed by a higher level reseller or service provider.
- Resale tariffs – are tariffs that the reseller charges their subresellers.
- Sales tariffs – are tariffs that the reseller charges their end users.

Wholesale tariffs are assigned to a reseller during their creation. *Resale* tariffs are configured as follows: when a reseller creates his subreseller, he clones the existing *wholesale* tariffs by applying a profit markup (a percentage) to them. Mapping between reseller's wholesale and resale tariffs is shown on

the **Usage Charges** tab. Sales tariffs are described in the **Product management** section below.



The screenshot shows the 'Add Subreseller' form with the following details:

- Customer ID: Pramod Kumar
- Customer Class: ABC Shuttle
- Currency: USD - US Dollar
- Permanent Credit Limit: 1000

The 'Usage Charges' tab is active, displaying a table with the following structure:

When I am charged using Wholesale Tariff	Charge Reseller / Subreseller using Resale Tariff	Description
Voice Calls		
Reseller - Outgoing	Clone the selected Wholesale Tariff	

Red annotations in the image indicate that step 1 is to select 'Clone the selected Wholesale Tariff' from the dropdown menu, and step 2 is to click the 'Update' button.

Resellers can always adjust their rates for *resale* tariffs for specific destinations. As a rule, rates in *resale* tariffs are higher than those defined for *wholesale* rates so that the reseller creates and receives a profit each time the service is used.

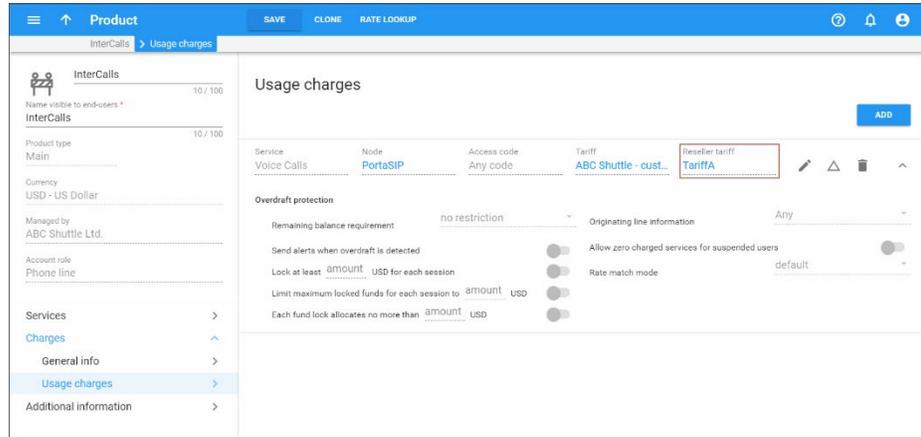
Handling a resale tariff change

When a service provider or reseller changes a resale tariff (for example, from tariff A to tariff B), the corresponding wholesale tariff for their subresellers also changes.

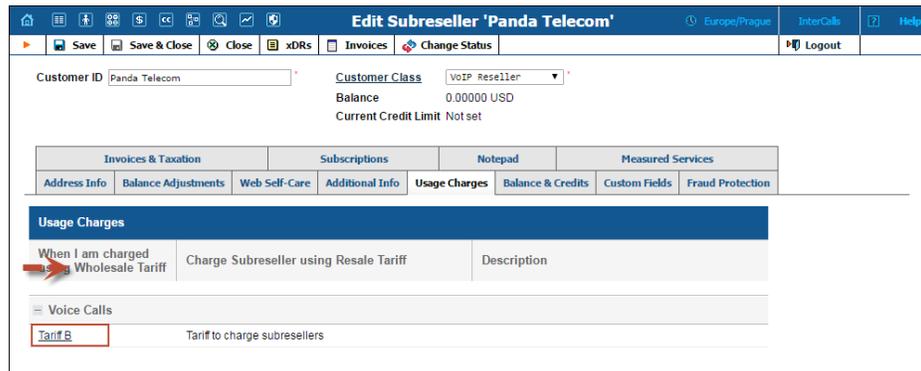
For example, the reseller network includes:

- The service provider,
- Company InterCalls, which is an authorized reseller for the service provider, and
- Panda Telecom, a local Chinese community that acts as a subreseller for InterCalls.

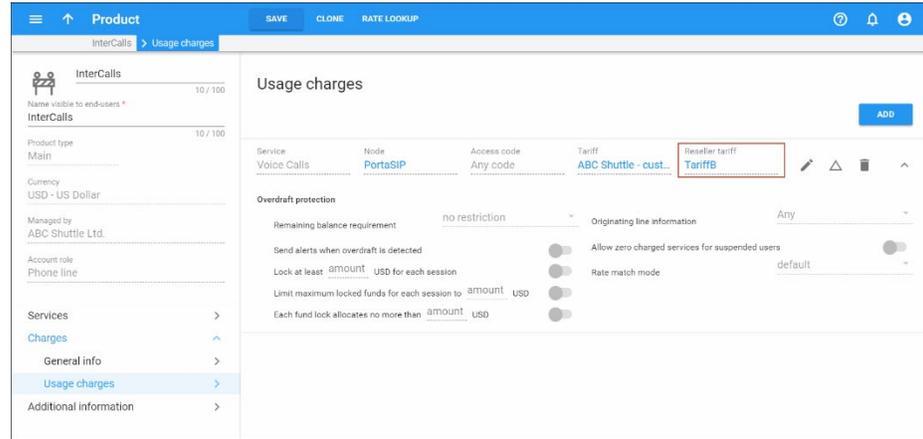
The service provider wants to charge InterCalls using tariff A. So, an administrator chooses tariff A as the *resale* tariff to be applied for InterCalls.



InterCalls can see tariff A as the *wholesale* tariff on their self-care interface (the tariff that InterCalls is billed by the service provider). InterCalls clones this *wholesale* tariff (by applying a profit markup) to create a *resale* tariff for charging Panda Telecom.



Later on, the service provider decides to charge InterCalls using tariff B. So the administrator changes *resale* tariff A to tariff B.

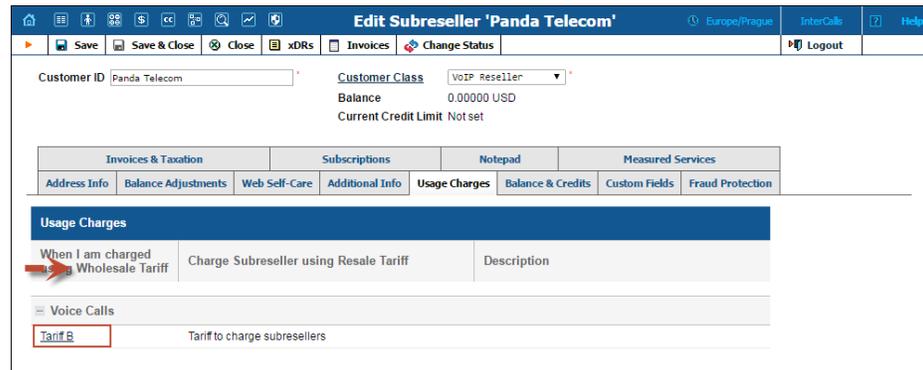


This tariff change results in the following:

- InterCalls can now see tariff B as their wholesale tariff.
- Tariff A is no longer used for charging InterCalls and thereby does not appear on InterCalls' self-care interface.
- The resale tariff that was created by cloning tariff A remains unchanged and is still used for charging Panda Telecom.

Be aware of the following limitation: calls made by Panda Telecom's customers that are in progress during this tariff change will not be billed. Therefore, we recommend that resale tariffs are changed only during off peak hours.

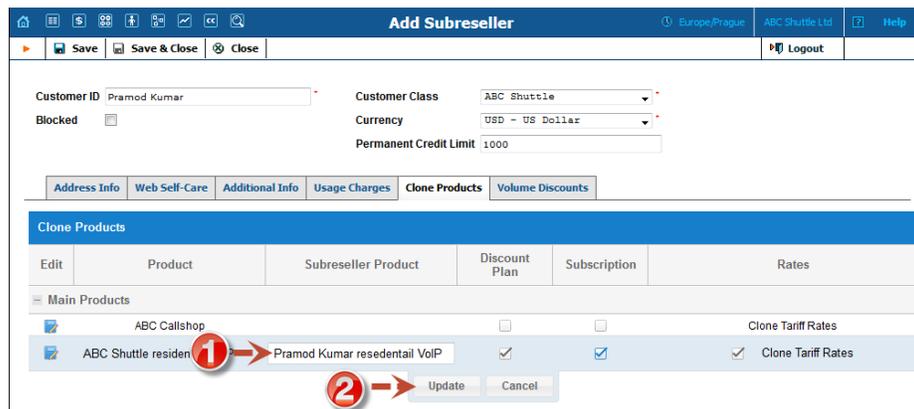
Alternatively, service providers and resellers can change the individual rates in their resale tariffs. This rate change does not affect the billing for calls that are currently in progress.



Product management

The product is the main tool that defines which services a reseller can offer to their end users. Resellers are not permitted to create their own

products by themselves as this requires deep knowledge of the ITSP's network structure (gateways, access lines, etc.), which most of them might not have. Instead, resellers are *assigned* their products from a service provider or higher-level resellers. This is done as follows: the higher-level reseller clones one of his products and the cloned product becomes 'at the disposal' of the lower-level reseller.



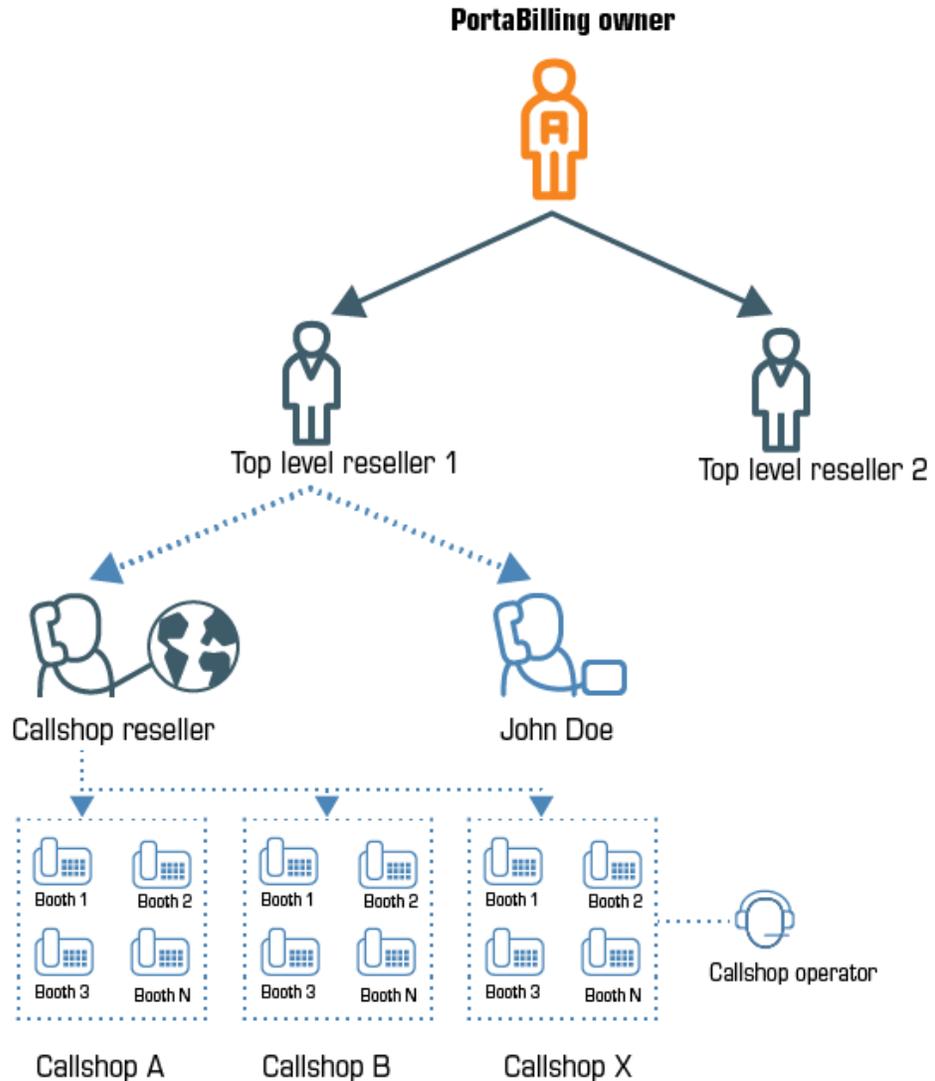
The screenshot shows the 'Add Subreseller' interface. At the top, there are navigation icons and the title 'Add Subreseller'. Below the title, there are buttons for 'Save', 'Save & Close', and 'Close'. The form contains several fields: 'Customer ID' (Prمود Kumar), 'Blocked' (checkbox), 'Customer Class' (ABC Shuttle), 'Currency' (USD - US Dollar), and 'Permanent Credit Limit' (1000). There are tabs for 'Address Info', 'Web Self-Care', 'Additional Info', 'Usage Charges', 'Clone Products', and 'Volume Discounts'. The 'Clone Products' tab is active, showing a table with columns: 'Edit', 'Product', 'Subreseller Product', 'Discount Plan', 'Subscription', and 'Rates'. The table lists 'Main Products' and a cloned product 'ABC Shuttle residen' (marked with a red circle '1') which is being cloned to 'Prمود Kumar resedentail VoIP'. The 'Clone Tariff Rates' checkbox is checked. Below the table, there are 'Update' and 'Cancel' buttons (marked with a red circle '2').

The cloned product is created with *sales* tariffs (i.e. tariffs according to which end users are charged). Sales tariffs, in turn, can come with or without sales rates, depending on whether the **Clone Tariff Rates** option was selected during the product cloning procedure. In the latter case, the reseller needs to define the appropriate rates for sales tariffs in order to sell the product.

With multi-level reseller functionality, ITSPs are able to easily and quickly build their own networks of resellers. This helps them increase their service offerings and gain higher profit margins.

Callshop resellers

A callshop is a public facility where end users come to make low-cost international or domestic calls. In a callshop, a customer is guided to enter a phone booth, makes calls and then pays for them at a cash desk. The callshop is managed by a callshop business owner who sets up the tariffs according to which end users are charged, pays for outgoing traffic and earns a profit on the margins between these costs. Callshop operators control the authorized use of booths, collect payments from customers and provide customers with receipts.



The following step-by-step scenario explains how a customer makes a call from a callshop and how he is then charged for it:

First, a customer is directed by a callshop operator to a specific booth.

1. The callshop operator unlocks a booth to allow making calls from this booth.
2. The customer dials a desired destination number.
3. PortaSwitch® identifies the phone booth, keeps track of call details, and calculates the corresponding charges.
4. When the call session ends, the customer leaves the booth.
5. The operator presents the customer with a bill for the call(s) made according to rates set up by the callshop business owner.

In PortaSwitch®, a callshop business owner is configured as a *callshop reseller*. This type of reseller *only* provides callshop services to end users (any other services are not permitted to be offered). In the multi-level reseller hierarchy, the callshop reseller is usually created under a top-level reseller or a subreseller.

NOTE: To permit the creation of callshop resellers directly under service providers the **Callshop** option must be enabled on the Configuration server.

The callshop reseller has their own self-care interface that has been specially customized for providing callshop services.



Equipped with a wide range of features that are oriented to callshop owners and their staff, the PortaSwitch® callshop solution assists callshop network owners to easily and effectively provide reliable, high-quality and cost-effective services to their customers.

Entities for common use by the administrator and resellers

To expand their service offerings to other spheres of business, ITSPs introduce the branded reseller program. In this program, the ITSP staff does all the service configuration and management while the resellers only deliver services to customers. This enables companies that lack deep technical knowledge to enter the telecommunications market and offer bundled solutions to customers under their name (e.g. a bank cooperates with an ITSP and resells a mobile application to their customer base).

To begin to offer the service, each reseller must have their own set of tariffs, products, subscriptions, etc. In case of a distributed reseller chain, this imposes an additional load on the administrator. Thus, to simplify reseller management, the administrator can share their own configurations with the resellers.

When the administrator marks the entity as shared, it becomes visible for all top-level resellers to use in their service provisioning. Since only the administrator manages the shared entity, the resellers see it as read-only.

The list of entities to share includes the following:

- **Products,**
- **Tariffs,**
- **Volume discount plans,**
- **Subscriptions,**
- **Customer classes,**

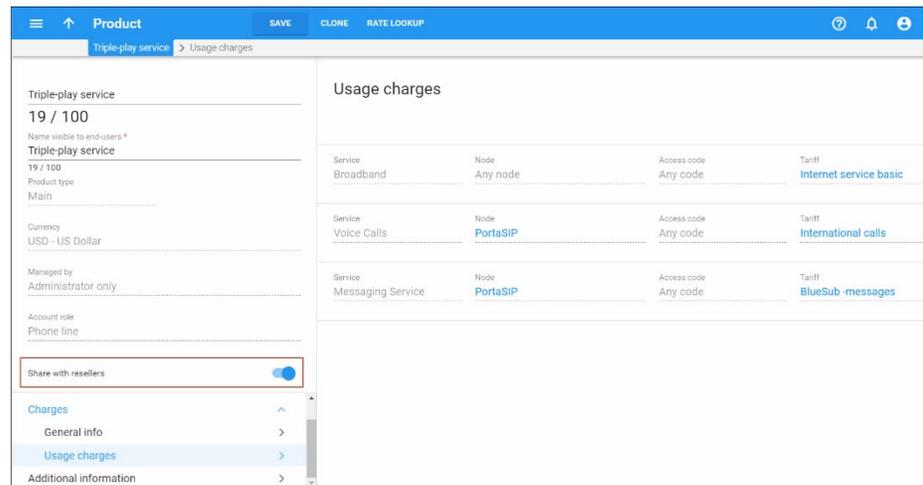
- CPE profiles,
- CPE.

Of course, a reseller can also create their own entities (e.g. subscriptions, CPE profiles, customer classes) and use them along with the shared ones.

Consider the following example:

The bank Openbank wishes to attract more premium customers. They introduce a Platinum service package that includes a personal bank officer, a premium credit card with a \$100,000 limit and even a travel SIM card for discounted service usage when roaming. Then they sign a contract with EasyCall to become their reseller.

EasyCall has several similar kinds of contracts with different companies. Therefore, EasyCall’s administrator, John Doe, configures the Triple-play product and marks it as **Shared with resellers**.



John Doe configures the reseller record for Openbank and grants them access to the system via the reseller self-care interface. Now Openbank has all the tools they need to sell services to customers and realize additional profit.

Summary of customer types

There are the following types of customers (all of whom have the same attributes, such as balance or address info):

- Reseller – Sells the service provided on your platform under his own name and product brand.
- Distributor – Resells your products to end users.
- Retail customer – Someone who actually uses the service.

In addition, all resellers can be divided into three groups:

- Top-level resellers – These resellers are created and managed by an administrator and do business directly with the service provider. A top-level reseller can create and manage his own subresellers and provide services directly to end users.
- Subresellers – These resellers are created and managed by a top-level reseller or another (higher-level) subreseller. A subreseller can create and manage their own (lower-level) subresellers and provide services to some end users.
- Callshop resellers – These resellers only provide callshop services to end users (any other services are not permitted to be offered).

Distributors are subdivided into two categories:

- Distributors – Distributors are created and managed either by administrators or by resellers and receive their commission from them. A distributor can sell products directly to end users or via its own subdistributors.
- Subdistributors – Subdistributors are created and managed by distributors and receive their commission from them even though it is the service provider's product that they sell. A subdistributor works with end users only.

Customers also can be subdivided into two categories:

- Direct customers – These are customers who communicate directly with the service provider (e.g., for account billing and payment). Direct customers can be resellers, distributors or retail customers.
- Subcustomers – These are retail customers and distributors which from a PortaBilling® standpoint belong to a reseller.

Retail customers, distributors and resellers are billed according to the credit model. Their balance reflects the amount they owe you. Thus it starts from a certain value (typically 0) and goes up with each call made, or down with each refund or payment. If the customer's balance reaches the maximum credit amount (credit limit), his accounts will not be allowed to use the service further. The only exception is a retail customer's debit accounts: since these are prepaid, and their calls do not affect the customer's balance, they can still make calls even if the customer has exceeded his credit limit.

Representatives

To win new customers and increase sales, service providers cooperate with various representatives (e.g. these people may sometimes be employees, independent contractors, sales agents, etc.). Representatives work on behalf of the company, selling services for a commission which is paid by the service provider.

Each sales representative is assigned initials. Their names and initials appear on sales forms, allowing you to associate specific sales representatives with specific sales in order to track their income.

This allows for:

- More accurate reporting of representative performance and commissions.
- Improved customer usage monitoring and analysis.

Commission calculation for representatives

PortaBilling® automatically calculates representatives' commissions. To make this happen, an administrator creates a commission plan wherein different calculation schemas are defined.

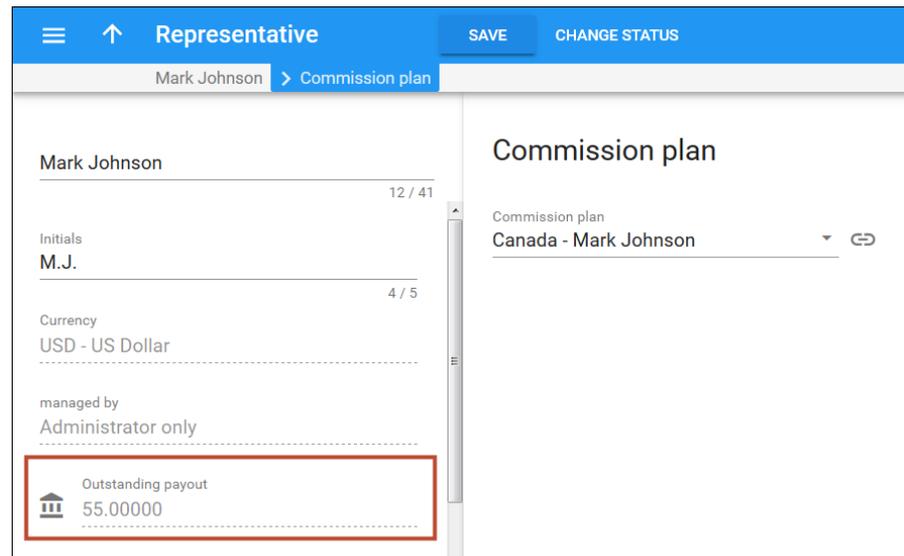
The commission amount depends on:

- service type (e.g. subscription, voice calls, etc.);
- commission type:
 - a one-time commission – a fixed amount of money that is paid one time only when a new subscription is assigned;
 - a percentage commission – a percentage of the customer's charges for services provided;
- the period – this is the number of months during which a predefined percentage commission is charged (e.g. for the first 3 months, the representative receives 10% of customer charges, for the next 3 months – 7% and for the following months – 5%).

There are two commission calculation modes that define when a commission is made available to a representative.

1. Payments.
PortaBilling® calculates and only makes available a commission once a customer makes a payment for their services. This mode is recommended for prepaid customers. Note that different commission amounts can be configured differently for time periods only. There is no option for differentiating a commission amount by service type and / or commission type.
2. Invoiced amount.
When a customer is charged, PortaBilling® calculates and makes available the representatives' commission. This mode permits flexible commission schemas to be configured – for example, to differentiate a commission by service type and apply a one-time commission for subscriptions. Note that a representative may want to receive their commission before the customer has actually paid.

When a commission is calculated, the representative's outstanding payout increases.



The screenshot displays the 'Representative' profile for Mark Johnson. The 'Commission plan' section is active, showing a dropdown menu set to 'Canada - Mark Johnson'. The 'Outstanding payout' field is highlighted with a red box and shows a value of 55.00000. Other fields include 'Initials' (M.J.), 'Currency' (USD - US Dollar), and 'managed by' (Administrator only).

Administrators and representatives can generate commission reports where all the commission details are presented (e.g. date, customer name, sum of charges, percentage of commission, etc.).

This is how it works:

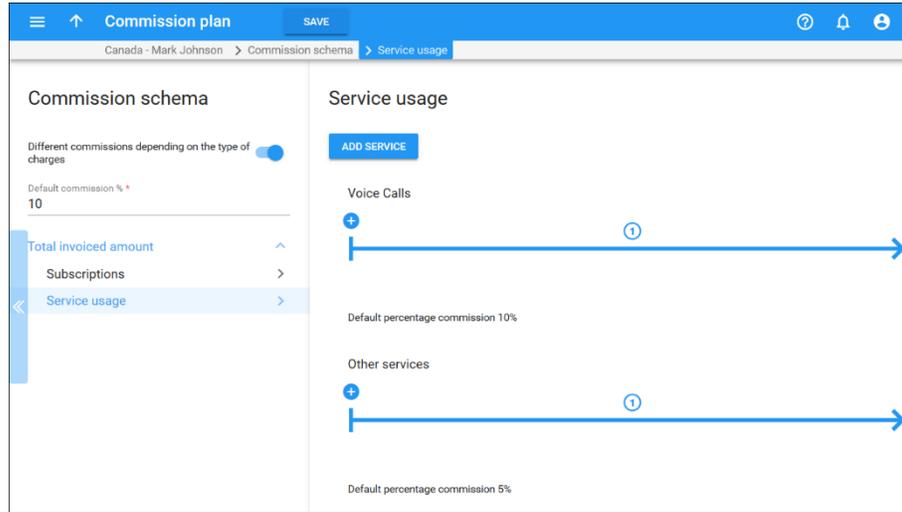
1. The administrator creates a commission plan for the new representative, Mark Johnson.
2. The administrator then assigns the commission plan to Mark.
3. When Mark brings in a new customer, the administrator identifies Mark Johnson as the customer's representative. Thus, Mark receives a commission from all of his customers according to his commission plan.
4. The administrator opens the representative's profile and sees the amount of money owed by the service provider to the representative.
5. The administrator therefore pays Mark his commission and manually decreases the service provider's debt.

Consider the following example.

Via Mark Johnson's commission plan, he will earn:

- a \$15 one-time commission – when the subscription “Equipment rental” is assigned.
- a 2% commission – for use of the \$10 monthly “Equipment rental.”
- a 10% commission – on all voice call service charges.

- and a 5% commission – on all other service charges.

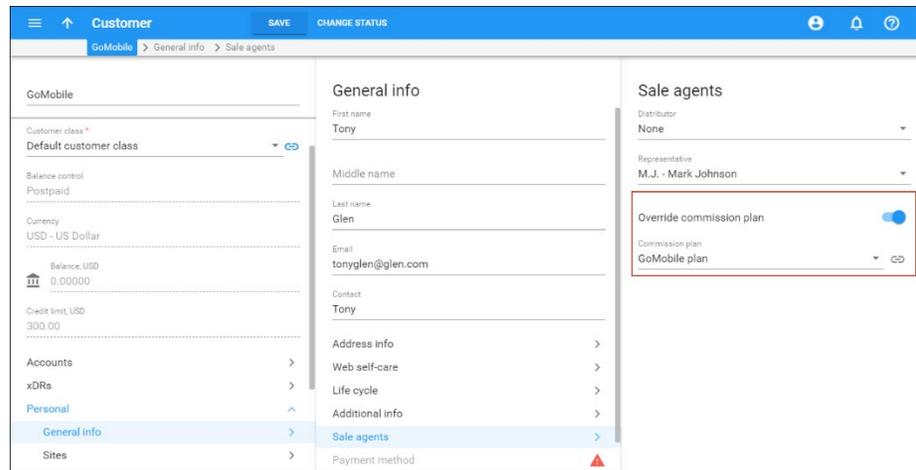


At the end of the month, the administrator opens Mark’s profile and sees that the total sum of Mark’s commissions comes to \$40.20. A report detailing the commission total is generated:

- \$15 one-time commission for the subscription;
- \$0.2 for subscription charge (2% x \$10 = \$0.2);
- \$15 for voice calls charges (10% x \$150 = \$15);
- \$10 for other service charges (5% x \$200 = \$10).

Override representative’s commission plan

It is possible that a service provider and a representative negotiate a different commission scheme for a particular customer. For instance, a representative receives a 5% commission according to his ABC commission plan but he wants a 10% commission for services consumed by the customer GoMobile. To make this happen, the administrator creates an individual commission plan and assigns it to that customer.



Currency conversion for representative's commission

It may happen that representatives and the customers they bring in operate in different currencies – e.g. USD and EUR. PortaBilling® automatically converts the currency for representative's commission.

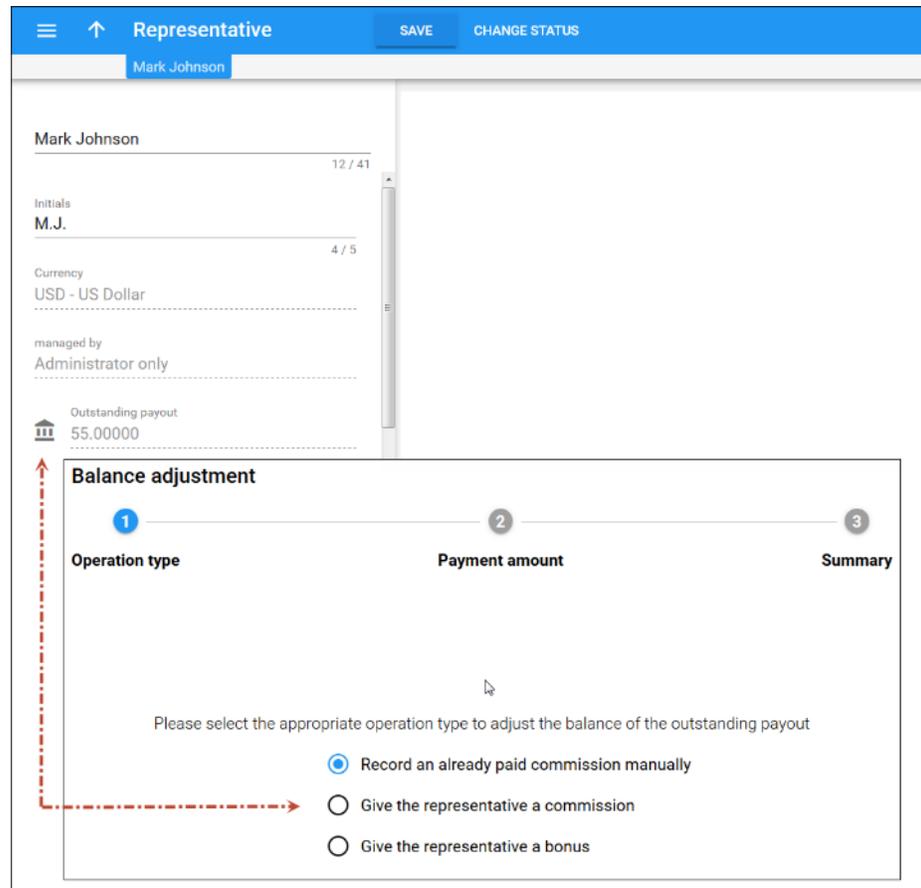
For instance, let's say your representative Mark Johnson operates in EUR. One day he earns a \$230 commission from his USD customer, so PortaBilling® converts the currency to EUR and Mark Johnson receives his €200 commission.

This allows to identify representatives for customers with various currencies and ensures that representatives' commission calculations and conversions are correct.

Commission payment

Commission amounts earned by representatives are shown in the **Outstanding payout** field. To adjust this sum, the administrator clicks  and the following options appear:

- Record an already paid commission manually – This option decreases the representative's commission amount when the service provider has already paid it.
- Give the representative a commission – This option allows to enter a commission amount for a particular customer and service and thereby increase a representative's commission amount.
- Give the representative a bonus – This rewards and encourages representatives: an administrator can issue bonuses – sums of money given in addition to commissions. This increases a representative's payout amount.



Representative: Mark Johnson

Initials: M.J.

Currency: USD - US Dollar

managed by: Administrator only

Outstanding payout: 55.00000

Balance adjustment

1. Operation type 2. Payment amount 3. Summary

Please select the appropriate operation type to adjust the balance of the outstanding payout

- Record an already paid commission manually
- Give the representative a commission
- Give the representative a bonus

This feature automates the commission calculation process for representatives. In addition, administrators can define different commission schemas by means of particular commission plans.

Vendors

Vendors are your service providers, for example termination partners or incoming toll-free line providers. Every time a call travels from your network to a vendor (via telephony or VoIP), a cost is associated with it. So at that point PortaBilling® will charge the account and customer for the call, and also calculate your termination costs.



For some services such as voice calls, vendors are required to provide consistent operation (e.g. call routing) and auditing (reporting, statistics) of your business. For other services (e.g. WiFi access or broadband Internet access) there are no direct costs associated with a particular session that are to be recorded as an individual xDR (for instance, your uplink provider may charge you a flat monthly fee for a 1Mbit/s channel to your WiFi hotspot), and so vendors will only be used to accumulate some usage statistics.

Connections to a vendor

Connections define points where calls travel from your network to a vendor. For connections via IP, we need to know the IP of the remote node (where the call is being sent to), while for telephony we need to know the gateway and, optionally, the name of the port the phone line is connected to. A connection defines a termination cost, i.e. the tariff according to which a termination partner charges you.

In case of a direct call between two users on your network, there is no “external” vendor. In this case you would create a special “internal” vendor which represents your own company. This will allow you to properly accumulate statistics for all on-net calls.

Connections from a vendor

It may happen in your business that you incur costs not only when terminating calls, but also when a carrier delivers calls designated for one of your customers to your network. Typical examples of situations in which such charges occur include the following:

1. You purchase a toll-free number from a carrier and interconnect with it via E1/T1. When a customer dials this number, the call is delivered to your gateway, and you can then provide a service (e.g. a prepaid card IVR) to the customer.
2. You purchase a set of phone numbers (DIDs) to be distributed to your customer's IP phones. When a number is dialed anywhere in the PSTN world, it will travel to the carrier who owns these numbers. The carrier will then forward the call either directly to your SIP server via an IP, or to your VoIP gateway via a PSTN trunk.

In both cases, you will be charged by the carrier who provides you with the service according to call duration, and these expenses must be reflected in PortaBilling® in order to show your cost/revenue figures correctly.

NOTE: Account and customer are billed **only** for outgoing connections. Connections from a vendor are purely for calculating costs.

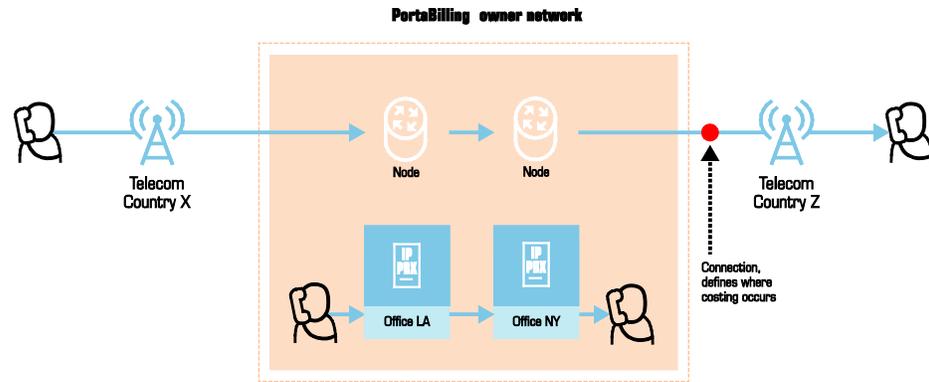


Figure 2-1

Connection load

PortaBilling® can build load graphs for every connection, so that you can monitor average load and quality parameters (e.g. ASR) for each one.

Nodes

Nodes are your gateways for facilitating calls, providing Internet access via WiFi, sending messages, or supplying any other type of service to the end user. Perhaps the most important feature of a node is that it can ask the billing for authentication or authorization and send accounting information to the billing system. It is very important that we exchange AAA data **only** with trusted gateways, also called **trusted nodes**. Usually, if this gateway is not owned by your company, but rather by your partner or customer, it will not be considered trusted, and so will not be created as a node in PortaBilling®.

Node ID, NAS IP address, and Radius source IP

What is the difference between these terms?

- **Radius source IP** address is the address that radius requests (UDP packets) come from. Radius will accept requests only from those IP addresses which are listed as **Radius source IP** for one of the nodes. This is used as protection against denial-of-service attacks and attempts to send fraudulent information to the billing system.
- **NAS-IP-Address** is the IP address of the gateway used for VoIP purposes. This address will be present in the radius data as the **NAS-IP-Address** attribute, and is used to identify which node sent the request. Usually **Radius source IP** is the same as **NAS-IP-Address**, except in a situation when your gateway has two

network interfaces, using one (internal) to communicate with the billing system and the other (external) for VoIP traffic.

- All of the incoming VoIP calls should be authenticated (see the **Implement Authentication For Incoming Voip Calls** section in the Unified PortaSwitch Handbook Collection for more details.). This also applies when a call comes from your gateway A to your gateway B. When node B (the terminating node) consults PortaBilling® as to whether the call should be allowed, it will send an authentication request to PortaBilling® with a User-Name containing some identification of the remote gateway. What information is used for identification depends on the application, but typically this is the IP address of the remote gateway. PortaBilling® will check if there is a node with such a **Node ID**, and, if so, the call will be allowed.

NOTE: You can only specify your own custom **Node ID** if you have enabled this advanced feature. Normally this is not required, and by default **Node ID** equals the **NAS-IP-Address**.

In most cases, all three parameters (NAS-IP-Address, Radius Source IP and Node IP) are the same.

Virtual environments

It is possible to have multiple isolated PortaBilling® instances, or virtual environments, within the same PortaBilling® installation. This functionality allows you to use PortaBilling® as an ASP (Application Service Provider) platform, or to create test environments which do not affect billing production activities.

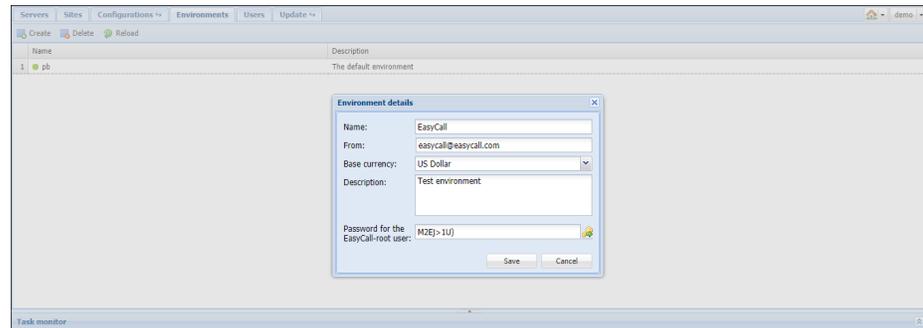
Each PortaBilling® virtual environment contains a complete set of data specific to that environment. A user who logs in to a user account for a specific environment will see only the users, customers, tariffs, accounts, xDRs, statistics, and so on for that environment. There are a few restrictions on the use of virtual environments to ensure their correct functioning:

- The login name for admin users must be unique across all environments. For example, you may not have a user ‘jack’ in both environment ABC-Production and environment ABC-Test. Rather, two users will have to be created, possibly ‘abc-jack’ and ‘test-jack’. The same applies to login names used by your customers / accounts / vendors on their self-care pages.
- A gateway (node) may be registered to only one environment. So you cannot have a node with the same IP address in two environments.

You may have multiple customer, vendor or account entries with the same name across environments. For example, you may have an account

'12345' in both 'ABC-Production' and 'ABC-Test' (from the previous example). In this case, the decision as to which account (from which environment) is determined by which gateway sends the accounting record (i.e. one registered in ABC-Production or in ABC-Test environment).

You may create new environments or delete them from the Configuration server web interface.



A new environment is created empty except for a single user account that you use to log in and then create other users, tariffs, customers, etc.

By default, the account's username is <env>-root, where <env> is the name of the environment you just created.

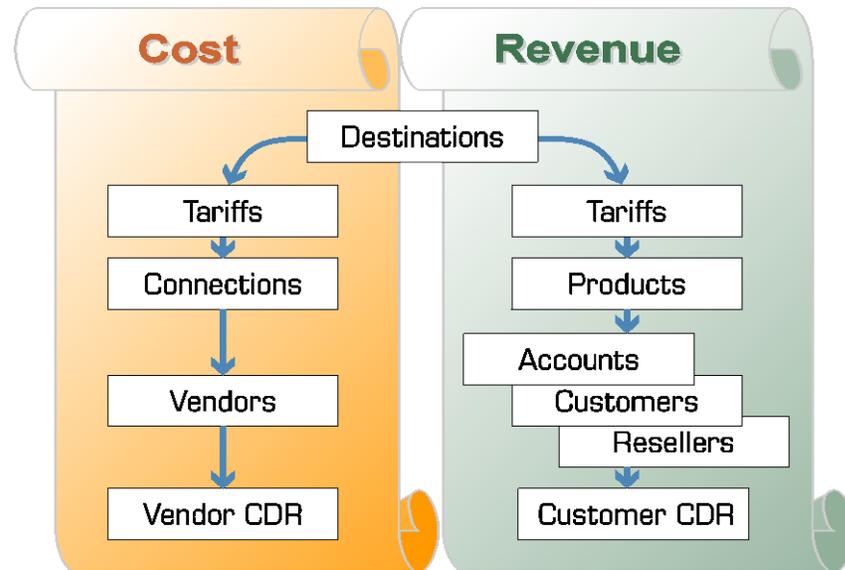
When you create a new billing environment, you define your email and base currency. Please note that you define the base currency only once and cannot change it once you save the environment. The default base currency is US dollars.

The user's password is randomly generated when you create the environment. So, if you create the environment EasyCall, to log into it you use the username EasyCall-root and then the generated password.

The password is displayed only during the environment creation step. Thus, you must copy the password before clicking the **Save** button.

Billing configuration

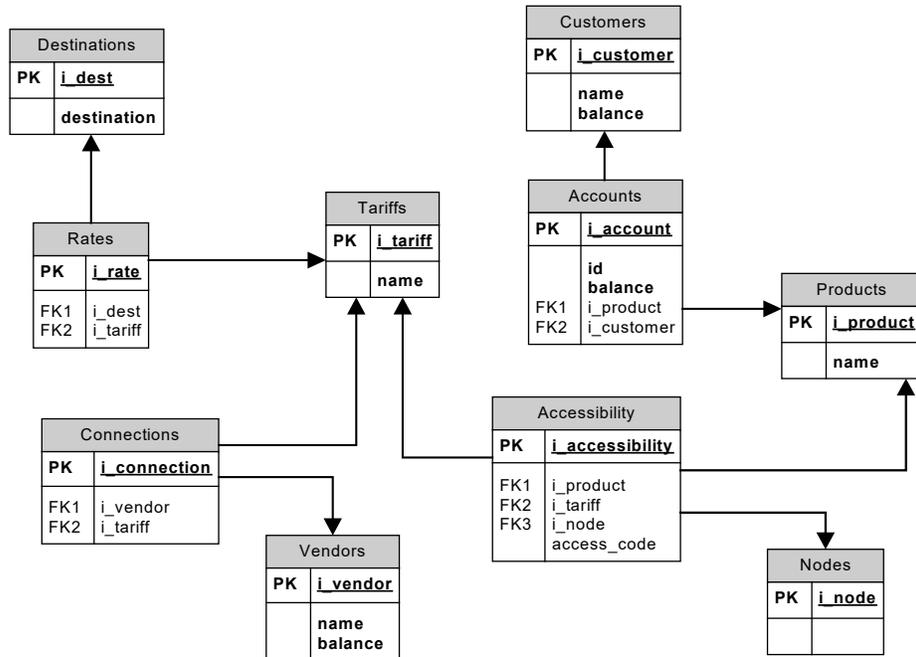
There are two independent parts of the billing configuration:



The Cost part provides you with information about expenses associated with running your VoIP business: call termination, incoming toll-free lines, etc. The Revenue part calculates charges applicable to your customers. Note, that these parts are independent (the only common item is Destinations). So, the amount charged to your customer is unrelated to your termination expenses and vice versa.

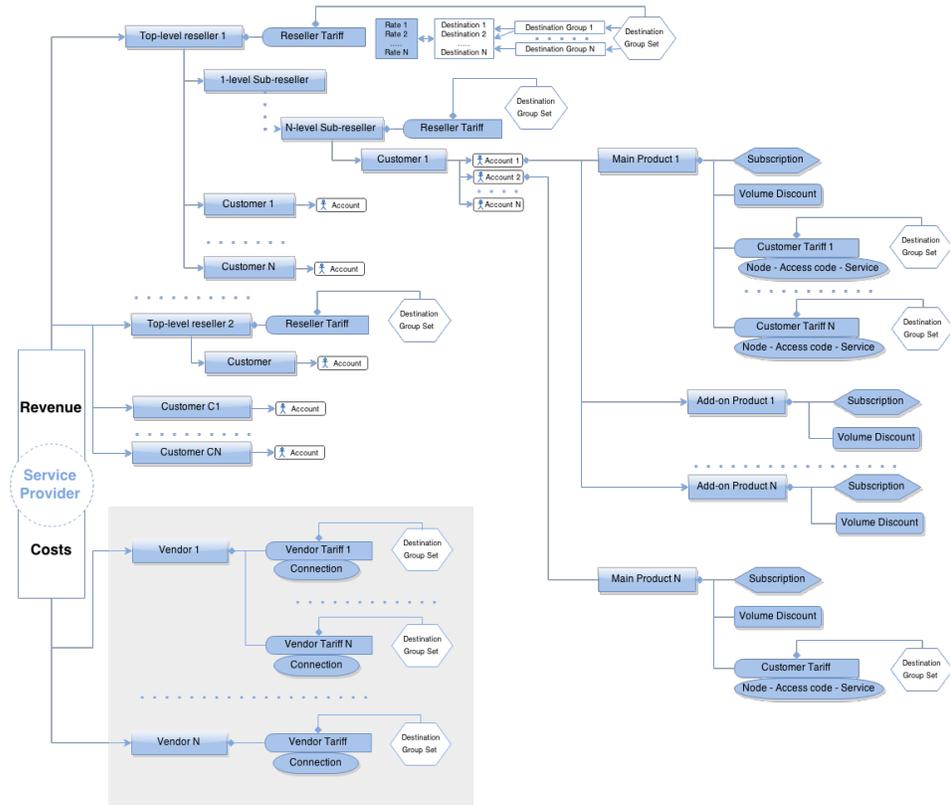
PortaBilling® data model concept

The illustration below shows some major elements and their interrelationships. For the sake of clarity, only the most important tables and columns are shown.



- Destinations define possible phone number prefixes which will be used in the system.
- Rates table stores information about rating parameters for a specific destination with a specific tariff.
- Tariffs serve as a pricelist to be assigned to a specific customer or vendor
- The relationship between vendors and tariffs is defined via the Connections table. A tariff assigned to a specific connection will be used to calculate termination costs for a given vendor.
- Every account is assigned a certain product, which defines how the account will use the service and how it will be charged for this.
- A product is related to tariffs via the rating table (Accessibility), which defines the mapping between the service consumption point (a combination of the node, access code and originating line information) and the applicable rating (service and tariff). Thus a single product may be linked with multiple tariffs.
- Accounts are linked to a customer, so that a single customer can have multiple accounts.

Billing entities hierarchy



The diagram above presents a hierarchical overview of the billing information stored in the system. On the cost side, under vendors, are located connections, with a tariff associated with each connection. On the revenue side, accounts are grouped under a customer (while a customer can, in turn, be placed under a reseller). The account is linked with the tariff (rate plan) via its product; note that this relationship is “one to many”, so that different tariffs can potentially be applied to the same account.

3 ■ Real-time charging

Destinations, rates and tariffs

Destinations, rates and tariffs are the essential parameters that define how a certain event must be billed. It is very important to find a correct set of billing parameters and rules (rates) for each event, and this is done by matching a service identifier with the rates available in the rate plan (tariff). The service identifier specifies exactly how the service is used. For example, the service identifier for voice calls would be the destination number (Called-Station-Id, or CLD, or DNIS). For SMS messages, it would be a combination of a mobile country code (MCC) and a mobile network code (MNC), while for services such as a pay-per-view movie it could be the movie category (“New Release,” “International” or “Classic”). Since rating telephony numbers is the most complex case, that will be the primary focus of this chapter.

Destinations

Destinations are a list of all possible phone number prefixes to be used in your system. You generally need a new phone prefix when you have a new service area for calls which are to be treated differently than others. For example, if you start providing calls to the Czech Republic, you should add destination 420 and specify it as the Czech Republic. Later, if you plan to charge calls to Prague differently than calls to the rest of the Czech Republic, you might need to add another destination with the phone number prefix 4202. All such destinations must be entered into the system before you can use these prefixes. This prevents errors and helps improve data quality. It is recommended that all of your phone number prefixes be defined in the E.164 format (i.e. <country code><area code><user’s local number> without the ‘+’ prefix).

For SMS delivery, however, you operate with mobile country codes (MCC) and mobile network codes (MNC) for phone numbers. Therefore, to charge users for an SMS to the Czech Republic via Vodafone, you would define the phone number prefix 42077 as the 230-03 code pair and add destination E.212-230-03 for the Czech Republic – Vodafone network in PortaBilling®. The format that defines phone numbers as a combination of MCC / MNC codes is called E.212 format.

It is virtually impossible to have an “ultimate” destination list that would contain all the prefixes for the whole world. First of all, it would be quite difficult to gather and maintain such information and second, and more importantly, having “all possible” prefixes would not offer any real benefits, and would only make rate maintenance more difficult.

For instance, if vendor A provides you with a rate of 0.10/min for calls to 420 (Czech Republic) and a 0.18 rate for calls to 420602 and 420603, you

must create three prefixes: 420, 420602 and 420603. If you also create the prefix 4205 (Czech Republic, Brno) it will only take up space in the database since it will never actually be used.

Therefore, we recommend that you follow a “required minimum” approach: only create those destinations for which you need specific rates for your carriers or customers.

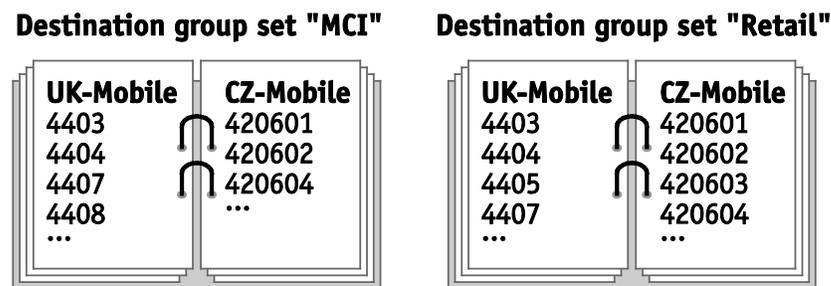
For service types other than those based on phone numbers, you may create symbolic destinations such as WIFI, NETACCESS or MOVIE-BLOCKBUSTER.

Destination groups

Sometimes you will have several prefixes for the same target destination, e.g. 420602, 420603 and 420737 for mobile numbers in the Czech Republic. All of these prefixes can be defined as the single destination group “CZ Mobile”, so when you enter a new rate for this destination group it will, in effect, create rates for each individual prefix.

Destination group sets

It happens quite often that different partners will assign a different meaning to the same destination groups. For instance, your partner A says that “CZ Mobile” is 420602, 420603 and 420737, but your partner B also considers prefix 420777 to be part of “CZ Mobile”. So you will need to have two versions of “CZ Mobile” – one for partner A and one for partner B. In this situation, you will create two separate destination group sets (“A” and “B”), as well as different destination groups inside those sets.



You can imagine each destination group set as a binder, with every page in that binder describing a certain destination group.

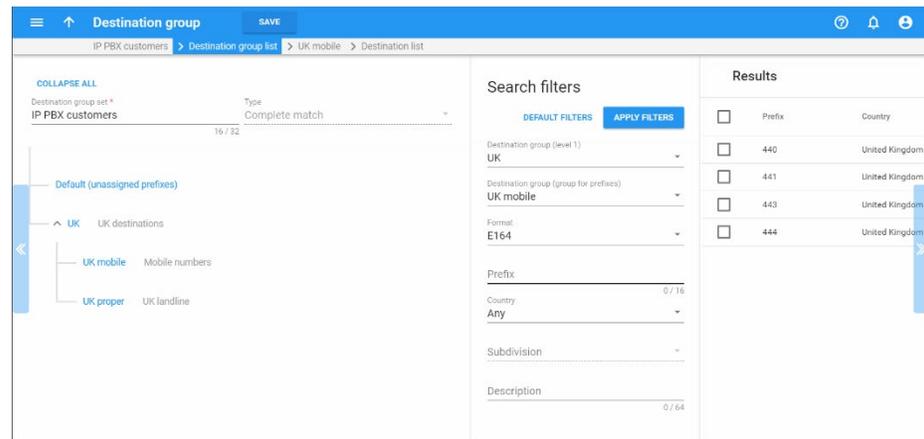
Complete destination group sets

In a complete destination group set any prefix that is used in the system belongs uniquely to only one destination group. This allows service

providers to clearly recognize traffic sent through their networks by their target destinations.

When an administrator creates a complete destination group set, it comes with a **Default Group** that includes all the prefixes available in the system. Upon adding new destination groups in the complete destination group set, the administrator moves the prefixes from the **Default Group** to these new groups.

A destination group in a complete destination group set (except **Default Group**) can have sub-groups. This allows the administrator to further distinguish the traffic flow. For example, to monitor calls sent to premium-rate numbers in the UK, the administrator defines **UK Proper** and **UK Mobile** sub-groups for the **UK** destination group. In this case, the UK destination group contains all of the prefixes added to its sub-groups.



At this time, complete destination group sets can only be used for configuring the fraud traffic profiles that monitor calling activities for the purpose of early toll fraud detection. However, in the future, complete destination group sets will be used for configuring volume discount plans, bundle promotions and tariffs.

Rates

A rate is a combination of billing parameters for a specific destination. For example, you can specify that calls to 420 are charged \$0.09/min during peak time and \$0.07/min during off-peak time, while calls to 420609 should not be allowed at all.

To easily manage the process of routine rate changes, each rate in PortaBilling® is assigned an exact time when it comes into effect (the **Effective From** attribute). This allows making an automated rate change in the future, so that a new rate can be entered today, but used only after

midnight on January 1st. When an existing rate is edited, this does not replace the actual rate, but simply creates a new rate effective from this moment, so that the PortaBilling® administrator can always browse the history of rate modifications.

Discontinuation of rates

Since it is important to have a clear overview of the history of rate changes, when a rate is not required anymore it cannot just be deleted from the system. For instance, you used to provide a special rate for the destination 442 London, UK, but as of March 1st there is just one rate for all of the United Kingdom. If a customer calls with a question about his bill for February, you should still be able to see why calls to London were rated in a particular way.

Therefore, rates which are no longer used are not “deleted”, but rather marked as “discontinued”. This has the same effect on call rating as if the rate was not there anymore. However, the administrator can still see all the history on the web interface, including the date when the rate became discontinued.

Tariffs

A tariff is a complete set of rates for a specific account, customer or vendor. Thus, it should include every possible destination allowed by their service plan (e.g. in the case of a telephony service, every destination to which you want to let them call).

It may be that the tariff will contain some prefixes that are part of more generic ones (for example, you will have a rate for the **420** prefix and for the **4202** prefix). In this case, the longest prefix match takes priority.

What happens if a destination is not included in the tariff, and the customer tries to call there? There are two possibilities:

- If outgoing calls are authorized via PortaBilling® (for example, by a Switching Server or by the prepaid card’s IVR) the customer will not be authorized to call this destination.
- If the gateway or switch which processes the call does not perform authorization in PortaBilling® before allowing the call to pass through, the customer will actually be able to make the call, but PortaBilling® will not be able to bill it properly, since the required information is missing. An email alert will be sent and a special CDR will be written into the database, so you will still have an overview of this call. It is highly recommended that you always use authorization of calls via PortaBilling®.

Wildcard destination

There is one special destination: | (pipe). When a rate for this destination is present in the tariff, it will match any dialed phone number. But this will only happen, of course, if there is no other, better match by the actual phone prefix. Such a rate should probably not be used in the customer's tariff, since it effectively authorizes him to call any phone number in the world. It is very useful, however, in tariffs for internal purposes. For instance, we can use it in the tariff used to produce CDRs for SIP-to-SIP on-net calls, since there the cost is zero regardless of which number is actually dialed. Otherwise, that tariff would have to include all phone prefixes for all possible phone numbers assigned to IP phones.

Also, for some services or billing features you may need to create other destinations. For example, in order to charge calls between extensions in the same IP Centrex environment according to a pre-defined rate, you will need the destination `VOICEONNETRX`. Consult the reference manual for a particular feature to find out which additional destinations it may require.

Relationship between destinations, rates and tariffs

Let's move from VoIP to a simpler scenario. Imagine you are the owner of an office supply store. You have to manage your inventory and price lists for customers and resellers.

- First of all, you must create a catalog of all the items you intend to sell. This is your internal document containing entries such as “20001 – ball-point pen, blue; 20002 – ball-point pen; black; 50345 – stapler;” and so on. Note that this is just a description of the items, without prices, since the price will depend on the specific vendor/customer.
- You will receive price lists from your suppliers and convert them into data files for every vendor, such as “Office Depot: 20001 – \$0.35; 20002 – \$0.35; etc.” Note two things: there is no need to include an item description in every file – you can always extract this from the catalog of items you have created above. Also, each data file will contain only some of the items, i.e. those which are provided by this particular supplier.
- The next step is to create similar price lists, but with the prices you will apply to your customers. Of course, you might have several such price lists: e.g. one for retail customers, another for business customers, yet another for resellers, and so on. Every data file will contain all of the items you offer to this category of customers, including prices.

Now let's come back to the VoIP business:

- *The catalog of items* corresponds to **destinations** in PortaBilling®. Every destination is similar to an item description in the catalog,

- i.e. it provides general information (e.g. 420 – Czech Republic proper, 420602 – Czech Republic mobile, and so on).
- *The price list (for vendor or customer)* is equal to a **tariff** in PortaBilling®. The price list includes all items the customer may purchase and all destinations the customer can make calls to. All tariffs are identical in structure, but some of them will be used to calculate your costs (vendor tariffs), while others will be used to charge your customers.
 - *A single line item in the price list* is equivalent to a **rate** in PortaBilling®. It gives the price per minute and other rating parameters for a specific destination.

Therefore, the standard sequence for setting up your service is as follows:

- Define the destinations you will use in your business. Basically you will need to define every unique prefix used by your vendors (an easy way to do this is by using PortaBilling® default destination set and PortaBilling® templates).
- Create a tariff for every vendor using the rate list they supply.
- Create a tariff for every customer / product. Once again, note that each tariff may contain a different set of destinations. For instance, the tariff for your vendor ABC may contain different rates for 420 (0.10/min), 4202 (0.09/min) and 420602 (0.18/min). But you can just list a single rate for your customer in the tariff, i.e. 420 – 0.20/min.

Relationship between tariffs and destination group sets

As mentioned before, multiple destination group sets may be defined in the system, with some destination groups (e.g. CZ-Mobile) defined for each of them. So, when you enter a new rate for CZ-Mobile in a specific tariff, which destination group definition should be used? To avoid such ambiguities, you will assign a destination group set for every tariff, which will be used to create new rates (if you do not assign a destination group set to the tariff, then you can only enter rates for individual prefixes).

Thus when you attempt to create a new rate for destination group CZ-Mobile, the following sequence of events takes place:

- PortaBilling® locates the definition of the CZ-Mobile destination group in the destination group set associated with this tariff.
- For every destination included in this destination group, the new rate is inserted.

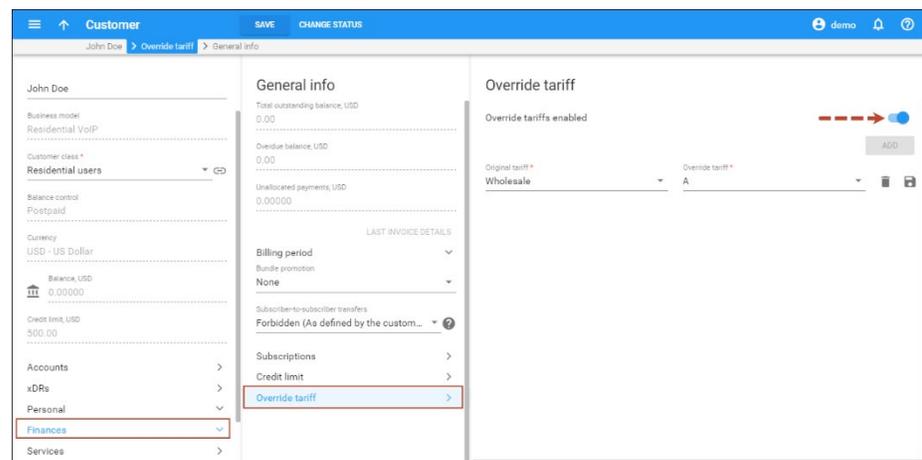
In other words, the result is the same as if you had created multiple rate entries manually. Rates which are created using a destination group do not differ in any way from rates created for a single destination. One important consequence of the foregoing is that if you change the destinations included in a destination group, it will not affect the previously created rates. Thus if you had 420602 and 420603 in the CZ-

Mobile destination group, and you now add 420737, this will not affect any of the tariffs. In order to have correct billing for the 420737 prefix, you must go to the corresponding tariff and add a new rate for the CZ-Mobile prefix.

Override tariff

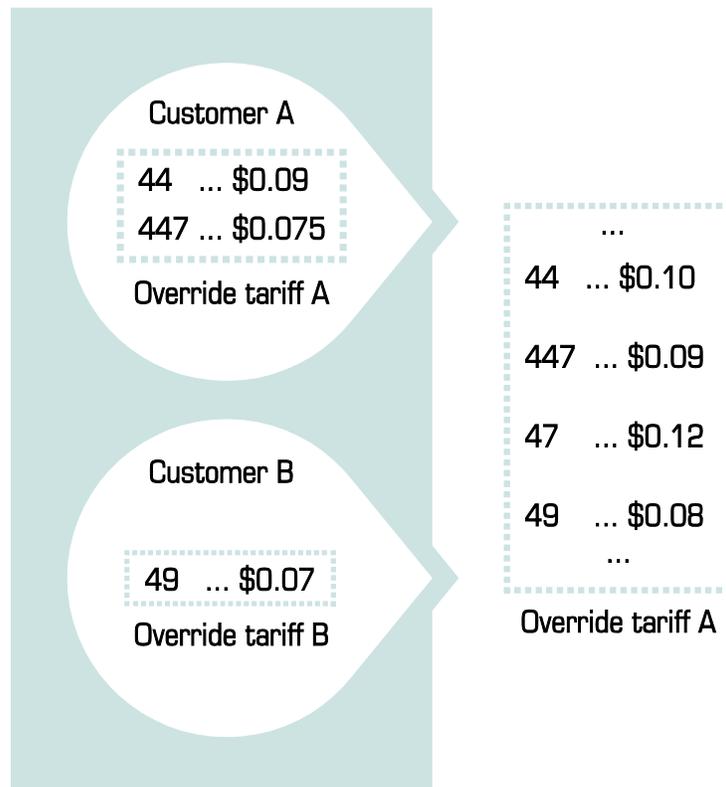
Often it happens that a particular customer negotiates an individual rate to a specific country, e.g. a company sending a lot of calls to UK mobile requests a reduced rate of \$0.09/minute compared to your standard rate of \$0.10, defined in the **Wholesale** tariff. A simple way to implement such preferential pricing is to create a volume discount for a specific destination group and assign it to the customer. The only disadvantage is that this discount is applied to the current price of UK mobile in the **Wholesale** tariff. So if the price there changes, the actual price for the customer will change as well. Of course, creating a dedicated tariff (containing rates for all destinations) just for this customer is also possible, but this significantly increases maintenance overhead.

In a situation where a customer requires a specific, pre-negotiated rate per minute for certain destinations, use the **Override Tariff** feature in PortaBilling®.



An override tariff is just like a normal tariff, but it contains only a handful of destinations, i.e. only those where the price must differ from the standard one. In our example above, it would be a tariff **A**, containing rates for 447, 448 and so on. The **Override Tariff** feature must be enabled in the configuration for a particular customer, and the administrator can then specify that the default tariff should be overridden by a specific tariff (in our case, the **Wholesale** tariff would be overridden by tariff **A**). Now, when calls arrive the billing engine will check for applicable rates in both the master tariff and the override tariff.

- If there is no match in the override tariff (i.e. there is no special price for this destination), then the rating is done as usual, using the rate in the master tariff.
- If a matching rate is found in the override tariff, and the master tariff does not contain a matching rate or the rate in the override tariff is more specific (the destination prefix is longer) than the rate in the master tariff, the rate from the override tariff is used. In our example, this would be \$0.09/min from tariff A.
- Finally, if there are matching rates in both the master and override tariffs, but the rate in the master tariff is more specific (the destination prefix length is longer), the override rate is ignored. This is done to prevent a “too generic” rate in the override tariff from erroneously granting a special price for some underlying and more expensive prefixes. In our situation, for instance, this means that if the master tariff contains a rate for 4489 (mobile premium numbers) at \$0.30, this is not a part of the special price arrangement, and the rate for 448 in the override tariff will not be applied.



Override tariffs allow the PortaBilling® administrator to easily create a custom rating for a specific customer (either adding new destinations not originally included in his tariff, or applying a special price for existing destinations). Also, it is possible to add multiple override definitions to the “override tariffs” table. This allows different override tariffs to be applied when a customer may be charged by several different master

tariffs (e.g. the customer can alternate between the “Premium” and “Cheap” routing plans).

Volume discount plans

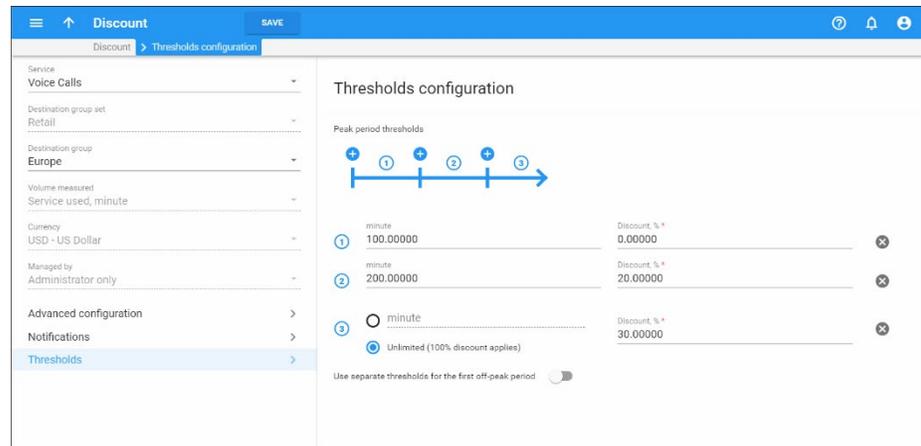
The volume discount plans tool is great for making automatic price adjustments and allotments. It allows you to adjust prices depending on the volume of services consumed by customers, allocate a certain volume of free services and transfer a balance to a sub-wallet that has been designed for a particular type of service. These are the three main discount schemes:

- **Discounts** – Using discounts, you can adjust prices depending on the volume of service consumed by customers. This encourages them to use the service more in order to receive the discount (e.g. spend \$20 for calls to the UK and get a 30% discount for the next calls during the month). You can also use discounts to make special offers (e.g. new customers receive 50% off for 10 GB of Internet traffic).
- **Quotas** – With quotas you are able to allocate a defined volume of services for your customers free of charge. A quota is usually combined with a recurring fee (a subscription). For example, provide 100 minutes for calls to Canada and 3 GB of Internet traffic for a \$30 monthly fee.
- **Service wallets** – This tool makes it possible to divide the customer’s balance into virtual sub-wallets. Each sub-wallet is designated for a specific service and destination group. Therefore, money transferred to a sub-wallet can only be used for a specific service (e.g. only for calls to the US or only for sending SMS, etc.). For example, a customer wants to have unbroken access to the Internet even if all his available funds have been consumed by making calls. So the customer tops up his Internet service sub-wallet by \$10 using a credit card. The customer receives 5 GB of Internet traffic that can be used even when his available funds reach zero. Service wallets can be topped up via the self-care or administrator web interface.

Use these discount schemes separately or combine them with each other to get the best solution for your business model. For instance, your service may include 300 free monthly minutes to US & Canada, 100 free monthly minutes to Western Europe, and a special offer of 15% off calls to India after the customer has called for more than 200 minutes. Each of the conditions above is represented by an individual discount scheme. Taken all together, they will be grouped as volume discount plan “EasyCall,” which can then be applied to a specific account.

Discounts

Use this discount scheme to provide a discount for a defined usage period – daily, weekly, semimonthly or monthly. When the usage period ends, the discount counters reset. For example, provide a monthly deal of the first 100 minutes for calls to the UK at a standard rate, the following 200 minutes at a 20% discount and all following calls to this destination at a 30% discount. When the next usage period begins, the customer’s calls to the UK are charged the standard rate.



You can also provide one-time discount. This discount type is a great tool for providing a permanent discount with no time limitation. For example, apply this discount type to provide a rate that is 10% cheaper for a wholesale partner.

A detailed description of how the billing engine calculates a discount is considered in the following example:

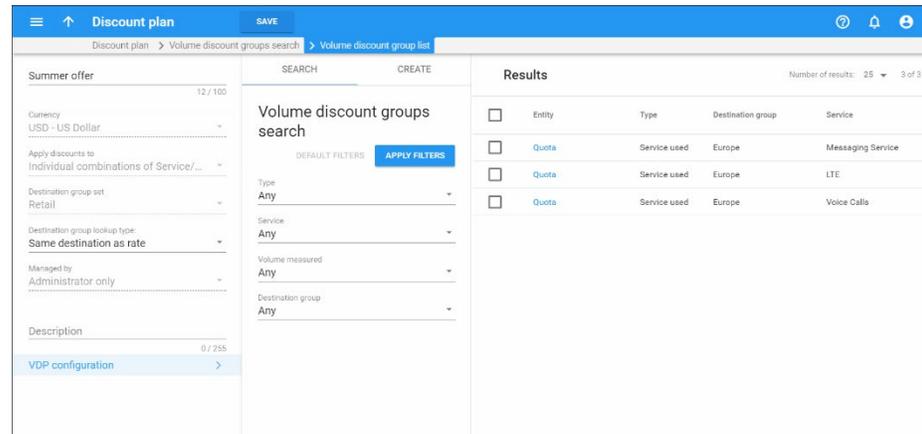
- You have the discount rule “first 200 minutes – regular rate \$0.20/min, 15% discount after 200 minutes” for calls to Israel.
- During one month the customer spends 230 minutes on the phone.

In this case, the customer will be charged $200 * 0.20 + (30 * 0.20 - 15\%) = 40 + 5.1 = 45.1$. Sometimes people get confused, and will assume that the charges in this case should be $230 * 0.20 - 15\% = 39.1$, which is **incorrect**. (In order to provide a discount to the total amount of services consumed use the **Bundle Promotions** feature).

Quotas

This is the most common way to provide a set of services for a recurring fixed fee. To do this, create a quota for your service, then create a subscription plan and assign them both to the same product. For example, you can provide a triple-play package that includes quotas such as 1000

minutes for domestic calls, 1 GB of Internet traffic and 150 domestic SMS – for a monthly subscription fee of \$40. Once the quota is consumed, the service is unavailable until the following month.



Entity	Type	Destination group	Service	
<input type="checkbox"/>	Entity			
<input type="checkbox"/>	Quota	Service used	Europe	Messaging Service
<input type="checkbox"/>	Quota	Service used	Europe	LTE
<input type="checkbox"/>	Quota	Service used	Europe	Voice Calls

Service wallets

Service wallet with zero initial balance

Sometimes customers want to reserve an amount of money for a specific service so they can freely use their available funds for other services. If a customer runs out of money at some point, the money in reserve remains available and the customer can still spend it only for the needed service. Within a single volume discount plan, the administrator can configure as many service wallets as needed. The volume discount plan can be assigned to a product, a customer or one of their accounts. When assigned, the service wallet with zero initial balance is empty, so in order to use it, the customer or administrator needs to top it up. Alternatively, there is the option of configuring a service wallet with an initial value.

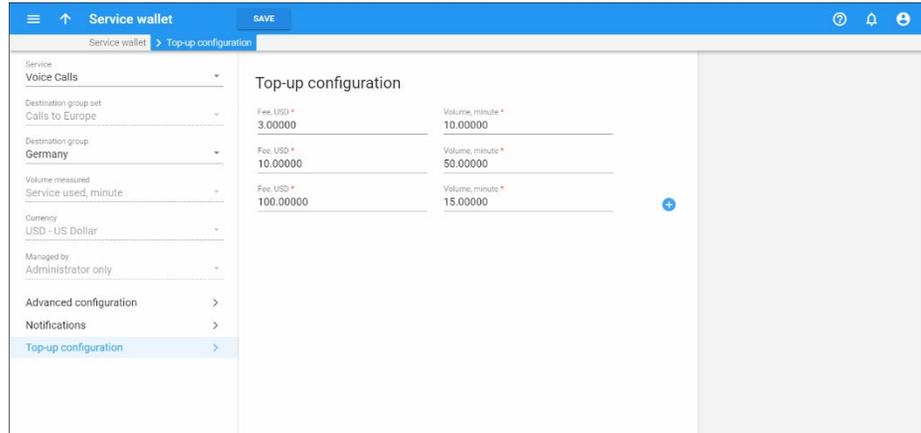
Service wallet with an initial balance

Use this service wallet scheme to grant a certain volume of service to your customers for one time. For example, grant your IPTV customers 1 GB of Internet traffic free of charge as a promotional offer. Once the traffic is consumed, the customer may decide whether to top up their service wallet for more Internet service or not.

Service wallet top up

This is a fast and easy-to-use tool to reserve customers' money for a specific service. The administrator or user can top up the service wallet at any time via their web interfaces. A top up is an efficient way to increase your payment flow since you, as the service provider, can sell a larger volume of service at a lower price. For example, you can offer 10-minute calls to Germany for \$3; 50-minute calls for \$10 and 100-minute calls for \$15. The user selects the most suitable offer among the three options and

either tops up his service wallet via the customer / account self-care web interface or requests this service from the administrator.

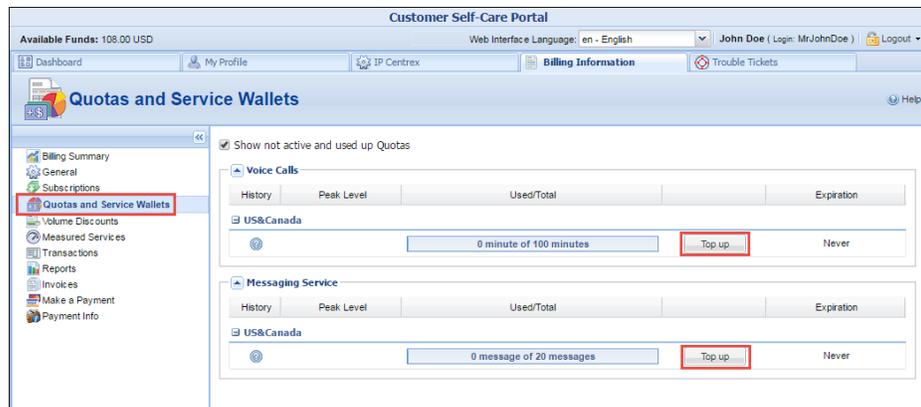


Consider the following example:

John Doe is a prepaid customer of EasyCall Inc. To be on the safe side, John Doe wants approximately 10 SMS and about 100 minutes for domestic calls in case he runs out of money using other services. He accesses the customer self-care interface and sees these available options:

- 20 SMS for \$3 and 50 SMS for \$6,
- 50 minutes for \$5, 100 minutes for \$9 and 200 minutes for \$15.

John Doe buys 20 SMS for \$3 and 100 minutes for \$9, his credit card is charged for \$12. Now he can use other available services (e.g. he surfs the Internet and makes international calls) feeling confident that this will not impact his ability to make domestic calls and send SMS messages.



Service wallet top ups using payment processors with external authorization

Customers can top up their service wallets by using payment processors that have external authorization such as PayPal, Moneybookers and PayArena.

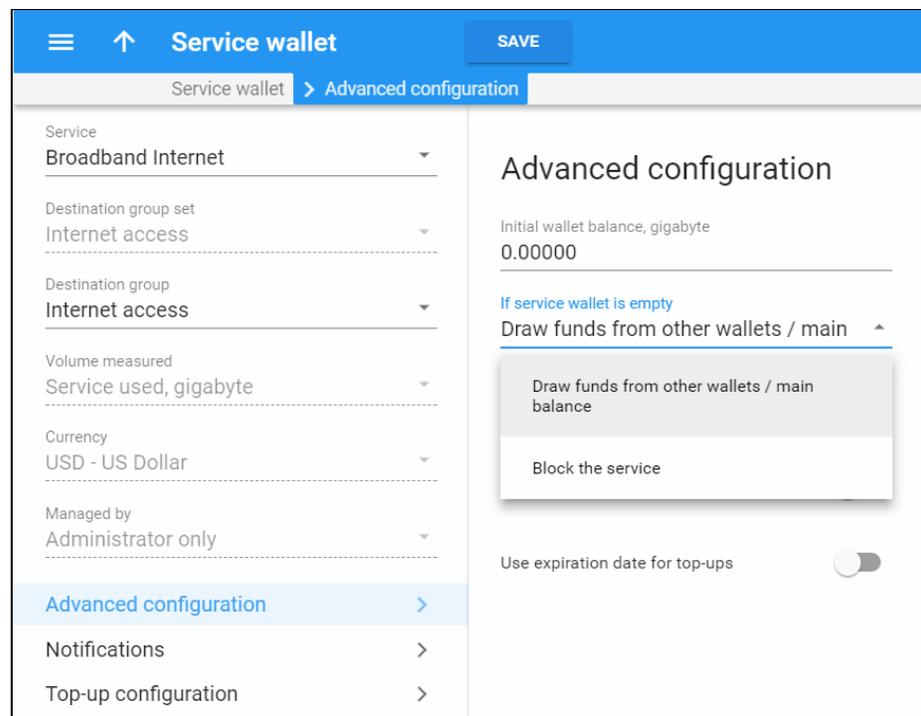
NOTE: The PayNearMe external payment processor does not support service wallet top-up transactions and therefore cannot be used for service wallet top ups.

When a customer initiates a top up from their self-care interface, they are redirected to the payment processor portal to provide their credit card information and confirm payment. Upon payment processing, the customer is redirected back to the self-care interface where they can view the updated service wallet’s balance.

Service restrictions for service wallets

With PortaBilling® you can fine-tune service provisioning for possible situations like when a customer’s service wallet for accessing the Internet is empty. For this situation, configure the system to:

- block service until the customer tops up their service wallet
- or
- continue to provide service but draw the funds from other applicable service wallets or from the customer’s main balance.



Thereby you clearly define the system’s behavior and also facilitate service provisioning.

Service wallets with a lifetime

One popular business model is to offer service provisioning for a certain period of time. This way, service providers encourage customers to make frequent top-ups so they can keep using the service.

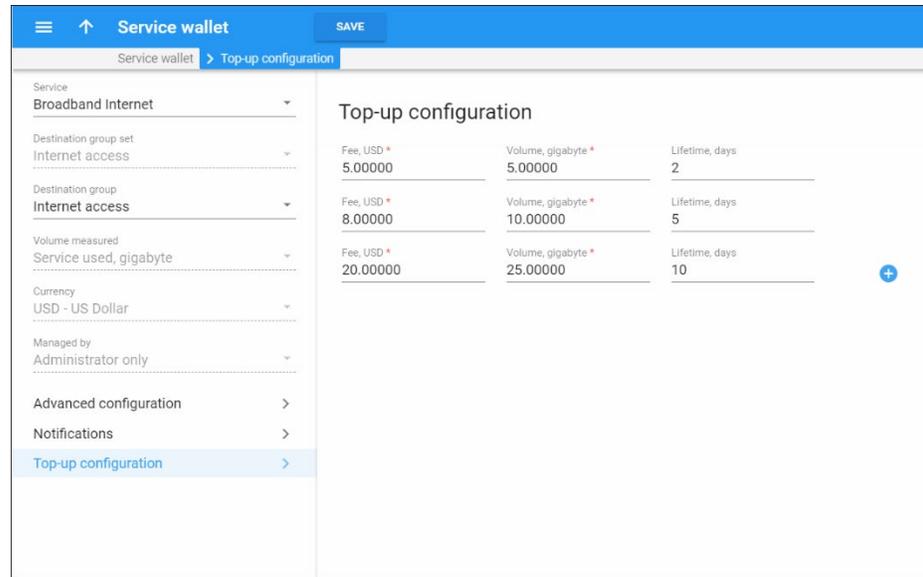
With PortaBilling®, you can introduce this business strategy by using service wallets that have a lifetime. Each top-up extends a service wallet's lifetime. Once a user makes a top-up, they can use the service for the period defined (i.e. the lifetime), after which it expires. Thus, the more a user tops up, the longer they can use the service. Note that service wallets that have an initial balance defined do not have an initial expiration date (i.e. the initial balance can be used unrestrictedly. It obtains its lifetime from its top-up).

If a user tops up a service wallet before it expires and there is an unused quota / balance, it is added to the top-up quota / balance. The service wallet lifetime, however, is not added: if the top-up lifetime is for less time than the current one, the service wallet lifetime is not extended.

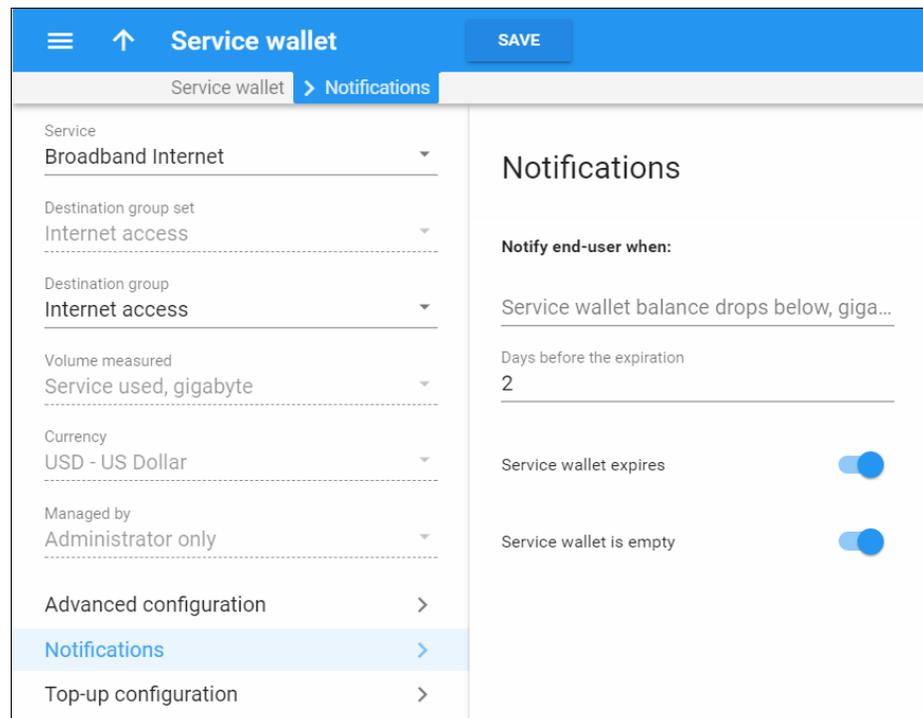
Consider the following example:

You are the Internet service provider and offer the following top-up options with your Start Internet service wallet:

- 5 Gb for \$5 and a lifetime of 2 days
- 10 Gb for \$8 and a lifetime of 5 days
- 25 Gb for \$20 and a lifetime of 10 days



John Doe has signed up for the Start Internet product and makes a \$5 top up. He sees that he has a 5 Gb Internet quota that is available for 2 days. During the day, he uses 4 Gb, so he needs to make another top up to continue surfing the Internet. He makes an \$8 top up and now he has an Internet quota of 11Gb (by adding the unused 1 Gb from the previous top up plus the new top up for 10 Gb) which is available for 5 days. If John Doe does not make another top up in 5 days, he will receive a notification about his quota expiration.



Note that if you grant a user some Internet quota (e.g. either as a promotion or compensation for quality issues), it will not affect the service wallet lifetime. Thus, if you provide John Doe with 1 Gb extra, his increased 12 Gb quota will still expire in 5 days.

By offering service wallets with different lifetimes, service providers can introduce marketing campaigns that facilitate service provisioning and boost sales.

Best practices

The table below matches the most common business cases and discount schemes suitable for them. Use this table to select the proper option for your own business model.

Description	Example	Scheme
The customer receives a permanent discount	The service provider grants a wholesale customer an unlimited 10% discount for calls to the US and Canada.	One-time discount
The customer has a monthly discount for a service	The service provider makes a special offer for business customers: “Spend \$100 for calls to Germany and get 50% off all following calls to this destination this month.” To receive the 50% discount again for the following month, the customer must spend \$100 for calls to Germany at the standard price.	Monthly discount
The customer receives a defined amount of service free of charge	The new subscribers of the “triple play” service receive 2 GB of Internet traffic for free.	Service wallet with an initial balance
The customer pays a monthly fee and receives a defined set of services	The customer pays \$30 monthly and receives: <ul style="list-style-type: none"> ○ free domestic calls ○ 30 free minutes for calls to Europe ○ 100 domestic SMS. 	Monthly quota (+ subscription)
The customer wants to maintain a separate balance for Internet service and for	The customer’s current balance is \$100. At the customer’s request, the administrator transfers \$20 from the main balance to the Internet service sub-wallet and \$10 to the international calls sub-	Service wallet with zero initial balance

international calls	wallet. The customer receives 5 GB of Internet traffic and 50 minutes for making calls, respectively.	
---------------------	---	--

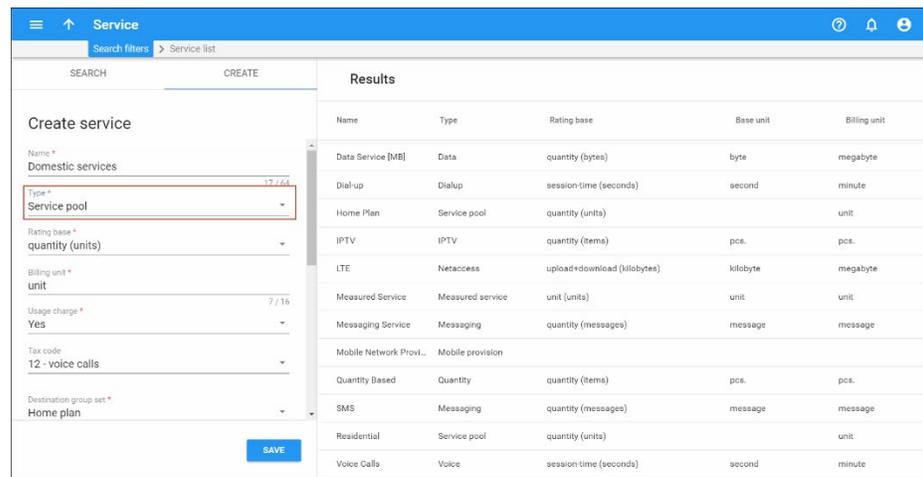
Quota / balance allocation for multiple services

Allocating separate balances for service usage in service wallets enables you to manage customer funds effectively. At the same time, your revenue increases, since customers top-up several service wallets at a time to access the service.

Another step in service usage management is to allocate a common balance that can be used for a subset of services. For example, allocate a \$50 service wallet for domestic calls, SMSs and Internet access and reserve the main balance to be used for international calling. Thereby you introduce quad-play services while monetizing their consumption efficiently.

To narrow down the services available within a common balance, consolidate them into **service pools**.

Every service within a pool can be measured differently (e.g. voice calls are measured in minutes, Internet access – in megabytes, etc.). Therefore, service pools are measured in unified virtual units.



Virtual units are used to express the cost for using a particular service that is included in the pool (e.g. one minute of calls costs 5 units while an SMS message costs 1 unit). If your pricing changes, you can conveniently re-configure it within the service pool.

Let's consider multi-service wallets and quotas, separately.

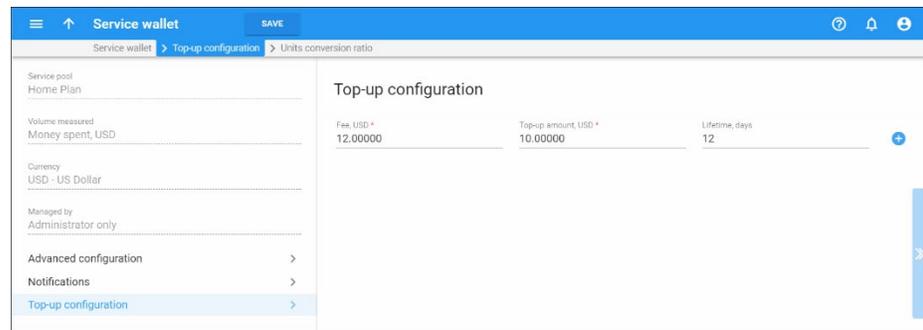
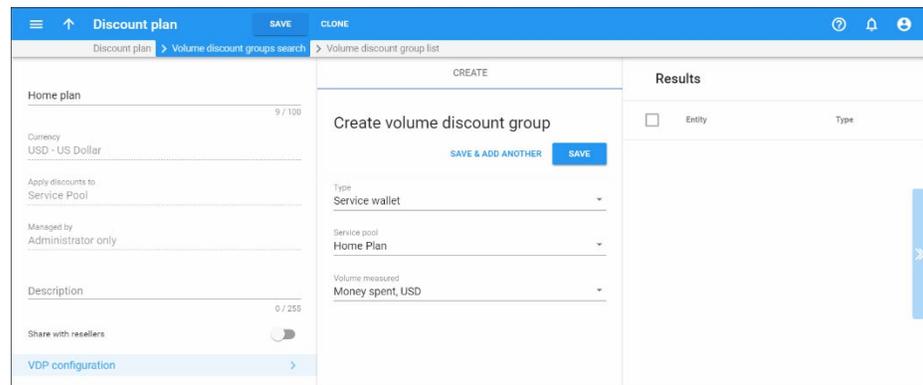
Multi-service wallets

Using multi-service wallets, you can allocate a common balance for your prepaid customers and allow them to use it for different services.

Similar to regular service wallets, in the multi-service wallets you can identify the service usage as being either monetary or in units. Encourage customers to use these wallets by introducing flexible top-up schemes (e.g. pay \$10 and use domestic services worth \$12 or pay \$20 and use domestic services worth \$25). Then when a customer tops-up their multi-service wallet, they can use it for any of the services included in the wallet while their main balance remains intact.

To illustrate how multi-service wallets function, consider the following:

You want to distinguish voice calls and domestic SMS delivery from international calls. To do this, create a “Home Plan” service wallet that includes Voice calls and Messaging services.



Prepaid customer John Doe lives in Canada and has \$5 worth of available funds. He signs up for the “Domestic Services” product and tops-up his service wallet with \$10. Each time John makes a call or sends SMSs to Canadian numbers, his \$10 decreases.

When John calls his mother in Australia, his \$5 worth of available funds decreases.

Quota consumption for multiple services

In some cases, you may not wish to reveal your pricing strategy to customers and / or competitors (e.g. prices are temporarily reduced to promote your triple-play service bundles and attract customers). Multi-service quotas work well for those cases. You offer customers quota consumption in virtual units and at the same time, flexibly define the pricing for each included service and re-configure it when needed.

Consider the following example:

Your “Paradise-100” quota provides 100 monthly units and includes services that are rated as follows:

- Voice calls to USA & Canada destinations – 3 units / minute;
- Voice calls to Europe – 10 units / minute;
- SMSs worldwide – 1 unit / message.

Customer Jane Doe signs up for this quota and starts using it. She calls her mother in New York for 10 minutes – so her quota decreases by 30 units, leaving 70 units available. After she makes a 5-minute call to a friend in the UK and sends 2 SMSs, she still has 18 units available for the month.

To fine-tune the service configuration further for your customers, combine the multi-service quota with regular quotas or service wallets.

For example, you provide domestic calls and SMSs within a multi-service quota while Internet access is provided from a separate service wallet. Thus, whenever a user makes a domestic call or sends an SMS, the quota counters adjust; and when they access the Internet, their service wallet balance decreases.

With this solution, service providers gain additional flexibility with their pricing and service management, and can be more competitive in the telecommunications market.

Customers, in turn, can form their own service usage patterns regarding on-demand consumption of units. In addition, topping-up a single service wallet saves time and increases customer loyalty to the service provider.

Mandatory parameters of discount rules

A discount rule is a set of parameters such as the service, thresholds, usage periods, etc. that defines how a discount scheme impacts billing for services. There are three mandatory parameters for each discount rule:

- **Service** – This parameter defines which service a discount, a quota or a service wallet is applied to. Select a desired service from the list that represents all available services in your billing

- environment. Note that within the same volume discount plan you can combine multiple discount rules for different services.
- **Destination group** – This parameter defines a destination group (a list of prefixes) that will be used in a volume discount definition. For example, if you wish to provide special rates for calls to Eastern Europe you can create a destination group that will contain prefixes for Russia, Ukraine, Belarus, etc. If a customer makes a phone call to 3806325551234 and a matching rate is defined in his tariff for the prefix 380632 (Ukraine, Mobile), the billing engine will search whether any discount rules in the volume discount plan are assigned to a destination group containing the prefix 380632.

NOTE: Within the same volume discount plan you can only select the destination groups that belong to the same destination group set.

- **Threshold type** – Discounts, quotas and service wallets can be configured based on the consumed volume of service (e.g. minutes, Mb, SMS) or on the service cost (note that this cost is calculated using a regular rate). For example, you can use thresholds based on call time (e.g. 10% discount after 200 minutes) or on the call cost (10% discount after more than \$10 worth of calls).

Phone number matching in volume discount plans

When volume discount plans are involved alongside “normal” rating as per the tariff plan, the billing engine must match the ID of the service used (e.g. the phone number dialed) with one of the destination groups used in the volume discount definition. The following volume discount plan lookup modes allow PortaBilling® to be fine-tuned for each individual case.

Same destination as rate

This requires exactly the same prefix as the one in the tariff in order to be used in the destination group for the volume discount definition. For instance, if a customer makes a phone call to 4206025551234 and a matching rate is defined in his tariff for 420602 (Czech Republic, Mobile), the billing engine will search whether any discounts in the volume discount plan are assigned to a destination group containing 420602.

By prefix or rate destination

Using this mode, the billing engine will look not for the exact same prefix, but rather for a destination group which contains a **matching** prefix. So if a call is made to 4206025551234 and there is a 420602 rate applied in the

tariff, the billing engine will search for a destination group in the volume discount plan that includes a prefix matching 420602; such matching prefixes could be 420602, 420, or even 4. In this case you could just enter 420 in your volume discount definition, and it will then be used for calls rated by different tariff rates (e.g. 4202 and 420602).

Therefore, with this mode you can avoid listing all applicable prefixes (420, 4202, 420602, etc.) in a destination group. However, such extensive lookups put a higher load on the billing engine, so they should be used with caution, and only in volume discount plans where they are truly needed.

By full rating pattern match

Another possibility is a full rating pattern match. This method disregards the destination used in the tariff entirely, and simply performs a lookup of applicable destinations in the volume discount plan using the phone number dialed by the customer. This is done in the same way as a rate lookup is done in a tariff; thus if the customer dialed 4206021345678, the billing engine will choose a discount with a destination group that includes 420, 420602, or any other destination matching the dialed number. Since several matches can potentially be found (e.g. a first discount is defined for destination group A, which includes 420, and a second discount for destination group B, which includes 420602), the longest match is used (in our example, this would be the second discount). As noted above, this is obviously the most resource-intensive option.

Usage periods

In order to provide more flexibility for creating new products, some of the discount schemes have their own “usage period” for each destination group. This period defines how often the volume discount counter is reset so that the customer starts with the first available discount level. This way the customer gets another pool of “free minutes” or, alternatively, the customer is returned to the standard rate and must send more calls through in order to obtain a discounted rate.

Typically you use the same discount usage periods as the customer’s billing cycle (so his free minutes “reset” when an invoice for the previous period is generated).

It is possible to apply usage periods that do not match billing periods: for example, 10 free minutes of calls to mobiles per *day* while the customer has a monthly billing cycle. Another common example is a “fair usage” policy for Internet access: although the user is not charged for the total amount of data transferred during a month, only 1GB of data is allowed to be downloaded per day, following which the bandwidth is severely limited.

Rollover of unused traffic

If at the end of the usage period (e.g. at the end of the month) there is unused traffic left (i.e. minutes, Internet traffic, messages, etc.), it can be rolled over to the next usage period.

For example, a customer has signed up for 100 bundled monthly minutes of free calls to Canada. By the end of the month, only 90 minutes have been used up. The 10 minutes remaining are rolled over to the next month, so during the next month 110 free minutes will be available for the customer.

The following example illustrates how rollovers work:

You create a discount rule with two rollovers and a monthly usage period. The unused traffic left over from the first month is rolled over to the second month and if not used completely, to the third month (two rollovers). If unused traffic is not completely used by the end of the third month, it expires.

If unused traffic from two or more usage periods is rolled over to the next one, the quota with the earliest expiration time is used first. Note that unused traffic cannot be rolled over more than three times.

Volume discount plan priorities

A volume discount plan can be applied to the following entities:

1. **Main product.** For better management, you can associate a volume discount plan with a certain product. This volume discount plan will then be applied by default to every account using this product.
2. **Add-on product.** There is an option for assigning a volume discount plan to an add-on product. When several add-on products with volume discount plans are assigned to an account, the volume discount plans are sorted according to the add-on products' priority.
3. **Customer.** Volume discount plans may also be assigned at the customer level. In this case, they would apply to the services consumed by any of the customer's accounts.
4. **Account.** You may directly assign an individual volume discount plan to a particular account.

If the account and the customer both have volume discount plans assigned for the same service and destination group and with the same threshold type, the discount plans are sorted according to their priorities. The priority is determined as follows (from the highest to the lowest):

- Volume discount plans assigned directly to an account.

- Volume discount plans assigned to an add-on product.
- Volume discount plans assigned to the main product.
- Volume discount plans assigned to a customer.

The volume discount plan with the higher priority in the list defines the way the next applicable volume discount plan with lower priority will combine with it.

The following example will help clarify how it works:

An account has a list of volume discount plans: one plan is directly assigned to an *account*, another is assigned via an *add-on product* and the last one is assigned via a *main product*. Thus, a volume discount plan assigned directly to an account takes the highest priority and defines how a volume discount plan added via an add-on product is combined with it. Then the volume discount plan that's assigned to the account via the add-on product, in turn, defines the way the volume discount plan that was added via the main product is combined with it. If any one volume discount plan applicable to a session is used up, the remaining volume discount plans are combined following the same rule.

Modes for combining discounts

There are several ways that discounts can be combined with each other. These are defined within the **Combine with other discounts** list and can be one of the following:

- Never
- Always
- When discount lower than 100% is active
- After reaching the last threshold limit.

Never

This mode of combining provides full override. It implies that a higher priority volume discount plan prevents the use of all lower priority ones. This means that even if the higher priority one is used up, the lower priority volume discount plans are ignored.

Always

Using this mode of combining, discounts are summed up to 100% (there is no money back for a consumed service). For example, two 30% discounts will result in a 60% final discount; 70% + 40% discounts will result in a 100% final discount; 100% + 30% discounts will result in a 100% final discount.

This mode of combining is also applicable to discounts that have different priorities. For example, the volume discount plan *Premium* provides a 20% discount for calls to the USA and is assigned to an account (the highest priority). The volume discount plan *Basic* provides a 10% discount for calls to the USA and is assigned to the customer (the lowest priority). To apply both discounts to the end user simultaneously, the volume discount plan *Premium* must have the combination mode **Always** applied to it since its priority is higher. If it does, the user can make calls to the USA using a 30% discount (20% + 10%).

In contrast to the **Never** mode of combining, the **Always** mode only applies to a volume discount plan that follows according to priority. For example, let's add another volume discount plan to the previous example. The *Standard* volume discount plan provides a 50% discount for calls to the USA, has the combination mode **Never** and is assigned to the add-on product. Thus, its priority is lower than *Premium*, but higher than *Basic*. Altogether, the user's discount will be 70% (20% + 50%), and the 10% discount will never be applied.

When discount lower than 100% is active

If a higher priority volume discount plan has a **threshold with a 100% discount defined within it**, only this volume discount plan is applied to a session until the 100% discount threshold is used up. Only after the volume of the consumed service exceeds the 100% discount threshold can the rest of the discounts be applied. If discounts defined in both the higher and the lower priority volume discount plans are applicable to the same call, they are summed up to 100%. For example, Germany discounted calls by 100% exclusively for 50-minute calls only are followed by a 50% discount for 1000-minute calls, while another discount for the EU offered a 30% discount. The EU discount would only become applicable if the 50 minutes are completely used, and then for a call to Germany, the total discount will be $30 + 50 = 80\%$.

After reaching the last threshold limit

A higher priority volume discount plan prevents the use of lower priority ones until its "unlimited" threshold is reached or the last (i.e. biggest) threshold is exceeded (regardless of discount value). For example, calls to Germany, discounted by 100% exclusively for 50 minutes are followed by a 50% discount for 1000 minute calls, while another (lower priority) volume discount plan for the EU offered a 30% discount. For calls to Germany the EU discount would only become applicable if the full 1050 minutes are used.

Let's have a look at the following example:

Your customer has account number 12126505550 in PortaBilling®. This account is assigned a *main product* that sets the rates (so that calls to the

US&Canada cost \$0.2/min) and two *add-on products* with different priority levels. Each add-on product includes a volume discount plan:

- A 50% discount on calls to the US for the first 60 minutes (add-on product with a high priority level; method of combining *after reaching the last threshold limit*).
- A 100% discount on calls to US&Canada for the first 20 minutes (add-on product with a medium priority level).

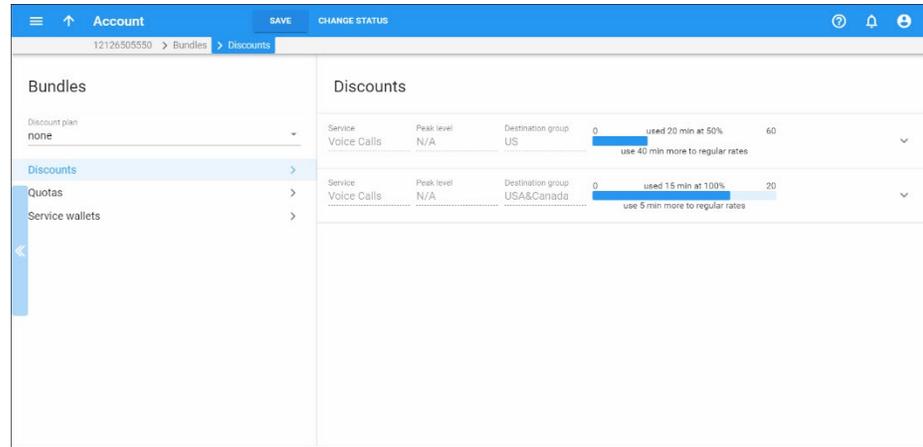
Service	Peak level	Destination group	0	used 0 min at 50%	60
Voice Calls	N/A	US			use 60 min more to regular rates
Voice Calls	N/A	USA&Canada		used 0 min at 100%	20
					use 20 min more to regular rates

The account number 12126505550 makes the next calls:

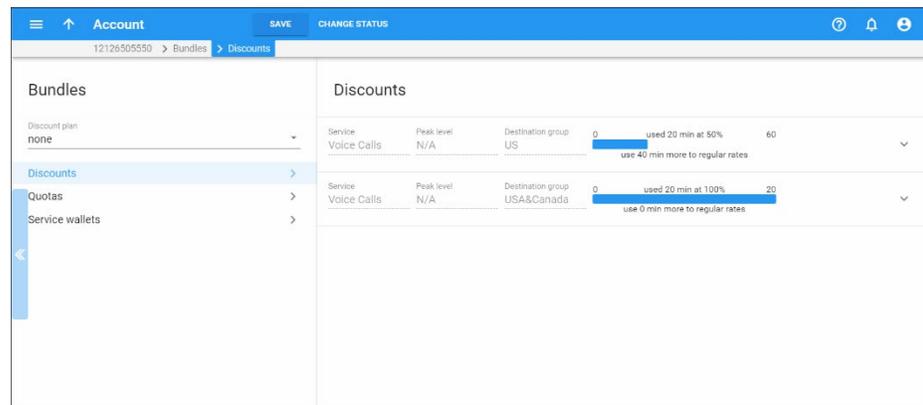
1. First, to Canada – for 15 minutes – it is charged using the “US&Canada” volume discount plan and comes to \$0.0.

Service	Peak level	Destination group	0	used 0 min at 50%	60
Voice Calls	N/A	US			use 60 min more to regular rates
Voice Calls	N/A	USA&Canada		used 15 min at 100%	20
					use 5 min more to regular rates

2. Then, to the US – for 20 minutes – it is charged using the “USA Cheap” volume discount plan (since it has a higher priority level) and comes to \$2 (20 * 0.2 * 50%).



- Finally, to Canada – for 10 minutes – it is charged using the “US&Canada” volume discount plan, but only for 5 minutes which comes to $5 \times 0 = \$0.0$, and then for 5 minutes without a volume discount plan which comes to $\$1 (5 \times 0.2)$.



This functionality allows the service provider to offer a flexible product portfolio that can be custom-tailored to market requirements and increase revenue via the sale of a la carte monthly calling plans.

Prorated thresholds in volume discount plans

If service is used during an incomplete billing period (e.g. customer signs up on April 20th and there are only 10 days left in the April billing cycle) then the volume discount thresholds can be prorated according to the duration of the period when service was actually used.

Let's consider a few usage examples:

- **Regular discounts** – A volume discount plan offers no discount for the first 100 minutes, a 10% discount for up to 200 minutes and a 20% discount after 200 minutes. The volume discount plan has a weekly usage period. If a user signs up for this plan on Wednesday then the first threshold is 71 minutes and the second threshold is 143 minutes since only 5 days remain in the current usage period. Thresholds are then 100 and 200 minutes for the next week and thereafter.
- **Change of volume discount plan** – Let's say Richard Roe switches over to a new product on November 15th that has a quota of 100 minutes. Richard receives 53 minutes ($100 \times 16 / 30 = 53$) of free calls for November.

Pay attention to how a volume discount plan is prorated when assigned on the last day of a current billing period:

- *before 10:59 p.m.* – services are available during the entire day. The quota is prorated, bringing the available amount of services for the whole day;
- *after 11:00 p.m.* – services are unavailable till midnight. The quota is prorated when the following billing period begins.

With this feature you can provide your customers with correct billing for incomplete volume discount usage periods, thus preventing your business from revenue losses.

Change of volume discount plans for customer / account

Preserving the counters

You can change a volume discount plan and preserve its counters for a new volume discount plan. This enables you to provide a single continuous offer to your end users with the option for adjusting / updating it at any time. To preserve the counters, ensure that both volume discount plans satisfy the following conditions:

- A destination group is identical and belongs to the same destination group set (e.g. both volume discount plans apply to the US & Canada destination group from the North America destination group set).
- Identical service (e.g. for both discounts the service is defined as Voice Calls).
- Identical threshold type (e.g. the provided discounts are based on the session cost).
- Identical usage period (e.g. both discounts have one-time usage or are reset monthly).
- The **Prorate thresholds for first usage period** option is enabled or disabled (e.g. when assigned in the middle of a billing period both discounts are recalculated depending on the number of days left until the end of the current billing period).
- Both volume discount plans must have the same priority level (e.g. the previous and new volume discount plans are associated with the main product or add-on products at the same priority level);
- The volume discount plan change must be done as a single operation. To do this, remove the old volume discount plan, assign a new one and click Save.

Note that when you change the volume discount plan via an add-on product, you can:

- Replace the already active add-on product with the new one. PortaBilling® applies the volume discount plan that's defined in the new add-on product at once.
- Schedule when the new add-on product will be activated. PortaBilling® will replace the active add-on product with the new one on the stated date.

If all conditions are met, the counters of the previous volume discount plan are preserved for use with the new one. Otherwise, PortaBilling® resets the counters. The following examples represent how changing the volume discount plans works when the counters are preserved:

Identical pre-condition

The service provider offers two bundles – 50 and 100 free domestic SMS messages with the identical pre-condition: valid after sending 25 SMS per month without a discount. A user subscribes to the 50 free SMS bundle. He has already sent 25 SMS without a discount plus 10 SMS out of 50 available and wants to upgrade to 100 free domestic SMS messages this month. He switches over to a new SMS bundle and since he has already sent 35 SMS messages, 90 (100-10) free domestic SMS are available for him immediately after the volume discount plan change.

New quota is smaller

A user uses up 9 GB and switches over to a new 7 GB product. He is not able to browse the Internet until the end of the month and starting from next month, he will only have a 7 GB Internet traffic quota available. The newly assigned volume discount plan will show the status **Used in Full** until it is reset.

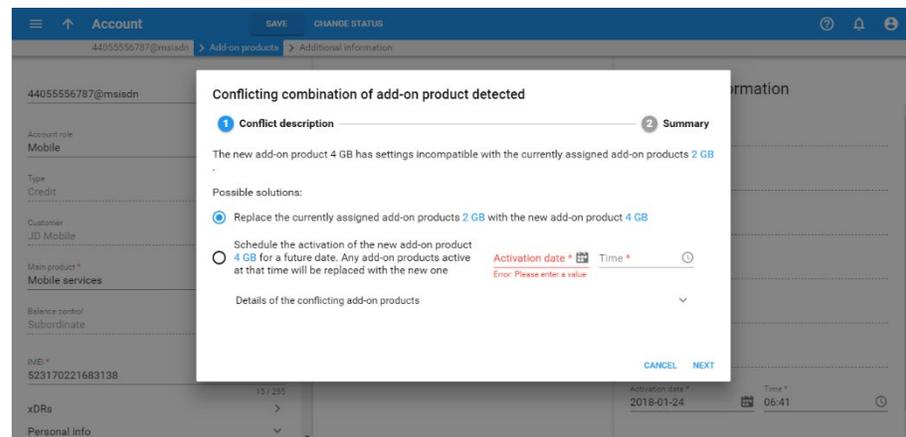
Single-step volume discount plan upgrade

Let's say that in addition to a Mobile Services product you offer both 2 GB and 4 GB Internet quotas via your add-on products. These add-on products have the same priority and are conflicting.

A customer who already has a 2 GB quota decides to call customer service and upgrade the quota to 4 GB.

When you assign an add-on product that has the same priority, PortaBilling® suggests that you:

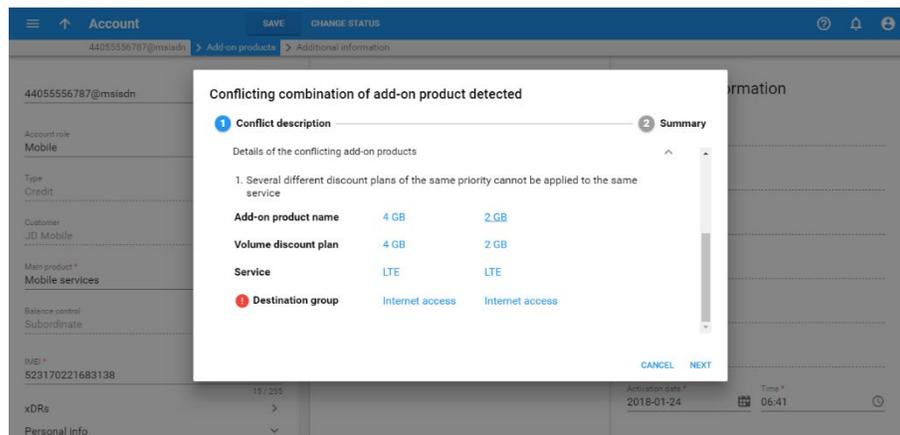
- Replace the already active add-on product with the new one. PortaBilling® applies the volume discount plan that's defined in the new add-on product.



- Schedule when the new add-on product will be activated. PortaBilling® will replace the active add-on product with the new one on the stated date.

Note that volume discount plan counters will be preserved during the product change.

You can see the reasons for the add-on product conflict. Just click the  arrow to the right of **Details of the add-on product conflict**.



Change of pro-rated volume discount plans

If a volume discount plan is changed / assigned *before* 10:59 p.m., the day of change / assignment is included in the proration for the newly assigned plan. Let's consider when a volume discount plan is:

- changed in the middle of a billing period; or
- assigned on the last day of a billing period.

Volume discount plan change in the middle of a billing period

Consider these examples when counters for the previous volume discount plan are preserved for use with the new one.

- Mary Smith is granted 60 minutes of free calls. By November 15th she wants to upgrade to a new quota of 100 minutes. She switches over to the new product and since she made 20 minutes' worth of calls during 14 days, there are still 8 minutes ($60 \times 14 / 30 - 20 = 8$) available from the previous quota. The newly assigned quota is prorated for the 16 days left till the end of the month ($100 \times 16 / 30 = 53$). With the unused minutes from the previous quota added, Mary receives a total of 61 minutes ($8 + 53 = 61$) of free calls for November.
- Mark Johnson does something similar. He made 59 minutes' worth of calls (almost the whole quota) during 14 days, so 31 minutes ($60 \times 14 / 30 - 59 = -31$) would be deducted from the new prorated quota ($100 \times 16 / 30 = 53$). Thus Mark only has 22 minutes available ($-31 + 53 = 22$) till the end of November.

Volume discount plan assignment on the last day of a billing period

If a volume discount plan is assigned on the last day of a current billing period:

- *before 10:59 p.m.* – services are available during the entire day. The quota is prorated, bringing the available amount of services for the whole day;
- *after 11:00 p.m.* – services are unavailable till midnight. The quota is prorated when the following billing period begins.

For instance, John Doe has a monthly billing period. At 6 p.m. on April 30th (the last day of his billing period) he assigns an add-on product that contains a quota of 100 minutes of free international calls. As a result, John receives 3 minutes of free calls ($100 \times 1 / 30 = 3$), since the quota was prorated for the last day.

NOTE: You can change both the threshold and the combination (**Combine with other Discounts**) option within the same discount and still transfer the volume of consumed discount. Please refer to the **Modes for Combining Discounts** section for detailed information about combination options.

Volume discount plan change using the API methods

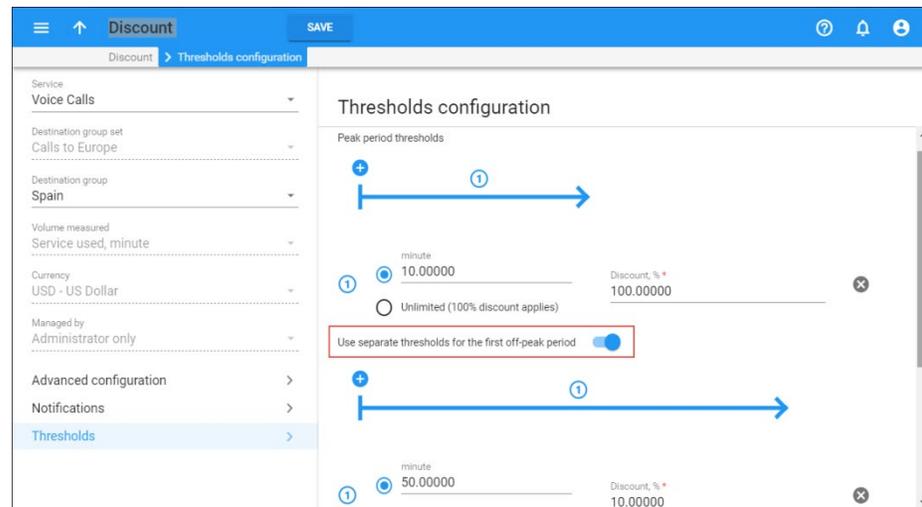
To change a volume discount plan using the API, use the following methods:

- To change volume discount plans included in the *main product* use the `i_vd_plan` property in the `Update_product` method;
- To change volume discount plans for a particular *customer* use the `i_vd_plan` property in the `Update_customer` method;
- To change volume discount plans associated with *add-on products* use the `assigned_addons` structure in the `Update_account` method. When using this method, the `assigned_addons` structure must contain the complete list of add-on products applicable to an account (i.e. the add-on products that will not be changed plus new add-on products that must be assigned to an account).

Please refer to the **PortaBilling XML / JSON API Reference Guide** for more information about API methods.

Peak and off-peak periods

By default, volume discount plans are applied irrespective of peak and off-peak periods, i.e. a minute is always counted as a minute. It is possible to create a volume discount plan with a different structure of discount thresholds for peak and off-peak time (and even for first and second off-peak periods). In this case, only counters for the relevant portion of the volume discount plan will be modified, based on the time when the service is used.



Threshold warnings

It is possible to configure the system so that a customer crossing a discount threshold (where a new discount level is applied) would automatically receive a notification. This could be used to notify the customer that he is either approaching the end of his discount provisioning or is now out of free minutes, or inform him that he is now eligible for a higher discount.

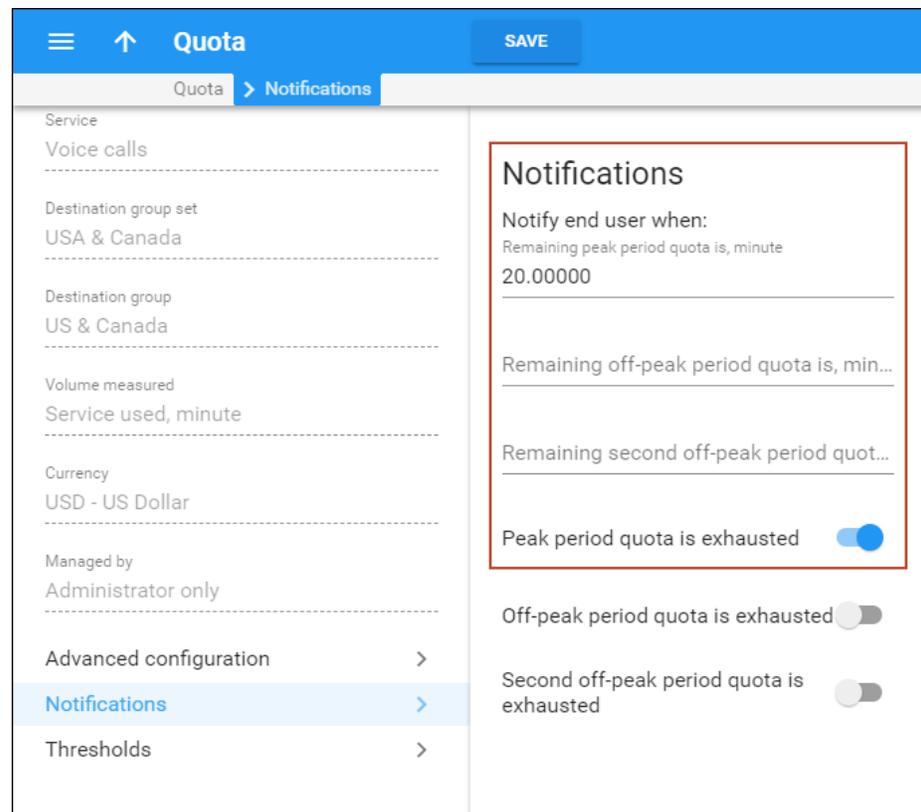
Customer notification before reaching a volume discount threshold

You can notify your customers that they are reaching the threshold for when a current discount stops being provided to them.

When a consumed volume of service reaches its warning threshold, as defined by an administrator for a particular discount, a corresponding notification is delivered to the customer by email or SMS.

Consider the following example:

A customer wants to sign up for 200 free minutes of calls to Canada. Moreover, he asks to be notified when he is down to only 20 free minutes. An administrator creates a volume discount plan for him that provides 200 minutes of free calls for Canadian destinations and then configures the requested customer notifications: to inform him when he has used up 180 minutes and again, when the whole 200 minutes are used up.



The screenshot shows the 'Quota' configuration page with a 'Notifications' sub-section highlighted. The left sidebar lists configuration options: Service (Voice calls), Destination group set (USA & Canada), Destination group (US & Canada), Volume measured (Service used, minute), Currency (USD - US Dollar), Managed by (Administrator only), Advanced configuration, Notifications (selected), and Thresholds. The main content area shows the 'Notifications' settings, which are enclosed in a red box. These settings include: 'Notify end user when:' with a value of '20.00000' for 'Remaining peak period quota is, minute'; 'Remaining off-peak period quota is, min...'; 'Remaining second off-peak period quot...'; and three toggle switches: 'Peak period quota is exhausted' (checked), 'Off-peak period quota is exhausted' (unchecked), and 'Second off-peak period quota is exhausted' (unchecked). A 'SAVE' button is visible in the top right corner.

Then he assigns this volume discount plan to the customer's account. When the customer has used up 180 of the 200 minutes for outgoing calls, he is notified by email accordingly.

NOTE: It is possible to specify warning thresholds only for value-defined quota or discount thresholds, i.e. notifications cannot be defined for unlimited calls within a volume discount plan.

Sending threshold notifications to your customers gives them a clear idea of the amount of services consumed and other charges. And it may also encourage them to buy additional discounted service bundles, thus raising your sales revenue.

Crossing discount thresholds and split xDRs

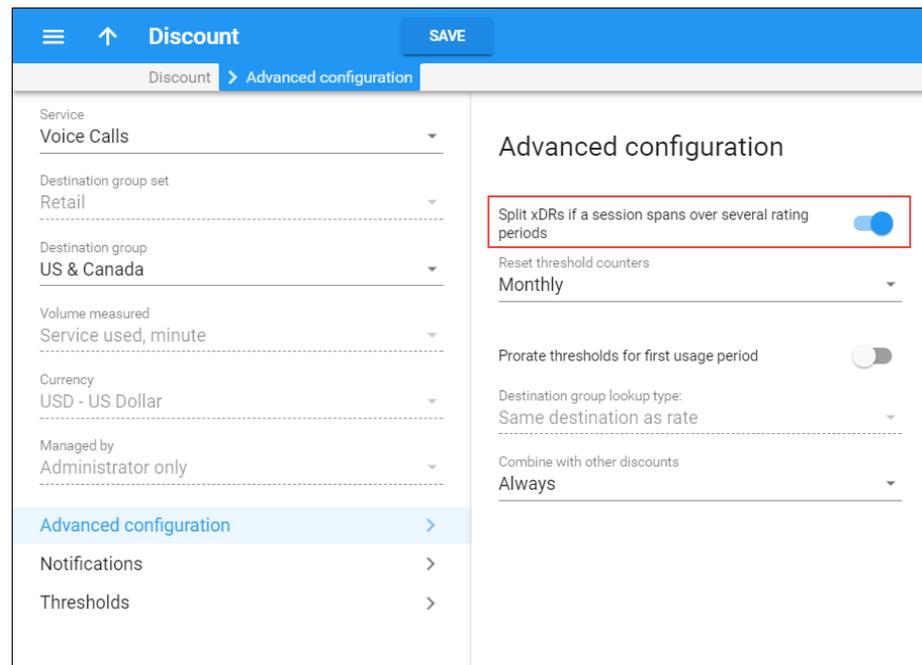
It may happen that a session (e.g. voice call or broadband Internet connection) spans several discount thresholds. For instance, a volume discount plan provides 100 free minutes of calls per month, after which the regular rate (\$0.10/min) is charged.

The customer has used up 98 minutes already, and then makes a phone call which lasts for 8 minutes. In this case, the first two minutes of this call are covered by the 100% discount and the normal rate applies to the remaining six minutes.

By default, an aggregated discount is calculated based on the rates/discounts applicable for each individual portion of the session and the duration of each portion. In our example the aggregated discount is: $100\% * (2/8) + 0\% * (6/8) = 25\%$. A call is then charged based on the total duration, the discount is applied and a single xDR is produced. In this case the amount charged for is: $8 \text{ minutes} * \$0.10 - 25\% = \0.60 .

A single xDR record makes it a bit more difficult to understand how the charges are calculated, but it correctly reflects the fact that the customer only used the service once.

If the **Split xDRs** option is activated for a particular discount level in the volume discount plan, then a session which crossed a threshold will produce multiple xDR records, each linked to the applicable discount level/rate.



In our example it is:

- xDR #1: $2 \text{ minutes} * \$0.10 - 100\% = \0.0
- xDR #2: $6 \text{ minutes} * \$0.10 - 0\% = \0.60

As you can see, the total amount charged is the same in both cases (\$0.60); it is just a question of which xDR presentation is better for your end users.

Transfer of funds and airtime among end users

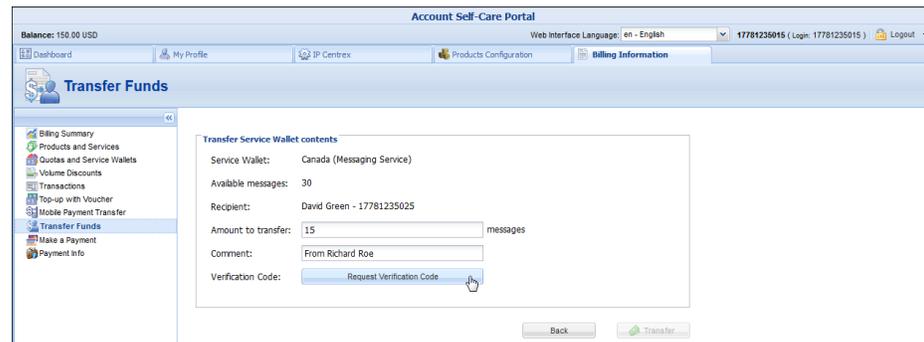
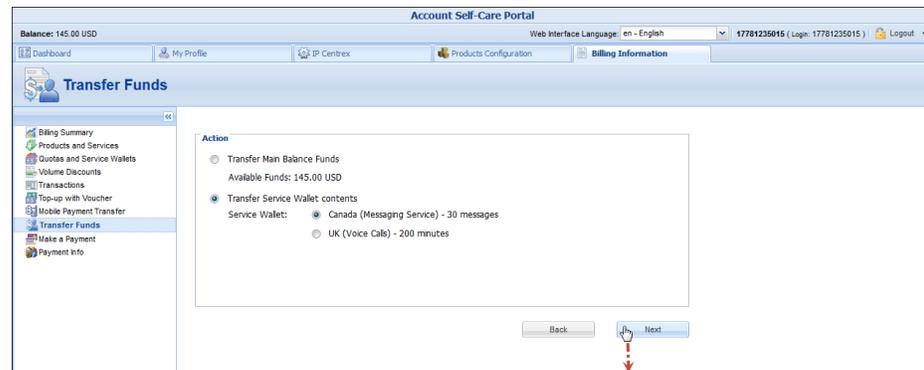
Transfer of funds and airtime among end users is a very popular service, especially in countries where payment systems are scarce. Money transfers provide a fast and convenient way to pay for everything without cash and / or credit cards.

With the transfer of funds and airtime feature, end users are able to:

- transfer their available funds to pay for goods and services, and
- share airtime from their service wallets with friends and family.

End users can take advantage of this service via their self-care interface. They can select what to transfer: either money or airtime (minutes, SMSs, GB, etc.) from their service wallet. Then they enter the recipient's phone number. For security purposes, each time a user initiates a transfer, the system sends an email or an SMS with a verification code. This code expires in 3 minutes and becomes invalid after the 3rd ineffective attempt to enter it. This is a safety provision to protect users from unauthorized persons obtaining access to verification codes.

As soon as an end user inputs a verification code to the system and the operation is successful, funds or airtime are transferred to the recipient.



Let's consider a transfer of funds and airtime, separately.

Transfer of funds

To prevent fraud, an administrator only enables the transfer of funds for an individual customer or a customer class.

End users who have credit accounts (with individual credit limits) or debit accounts can transfer their available funds from their self-care interfaces. This is very convenient, since it encourages them to manage their funds efficiently, and be ready to either top-up their relatives' balances or pay for goods in a shop.

For instance, let's say the user Richard Roe has \$50 worth of available funds. He decides to transfer \$35 to pay for goods in a grocery store. Richard Roe visits his self-care interface, enters the phone number of the grocery-store owner and the amount of money to be transferred – \$35 in this case. He receives an SMS with a verification code and immediately enters that into the system. The system accepts the code and funds are withdrawn from his account. The grocery-store owner receives \$35 and Richard's available funds are now \$15.

If an end user makes a transfer by mistake, an administrator can reverse the transaction.

Transfer of airtime

Transfer of airtime involves the transfer of the contents of any service wallet – minutes for calls, SMSs, GB of Internet traffic, etc. The transfer is only allowed among end users having identical volume discount plans. This is a safety measure to prevent service abuse (e.g. when end users sign up for cheap products and then transfer airtime to use for premium products or commercial purposes).

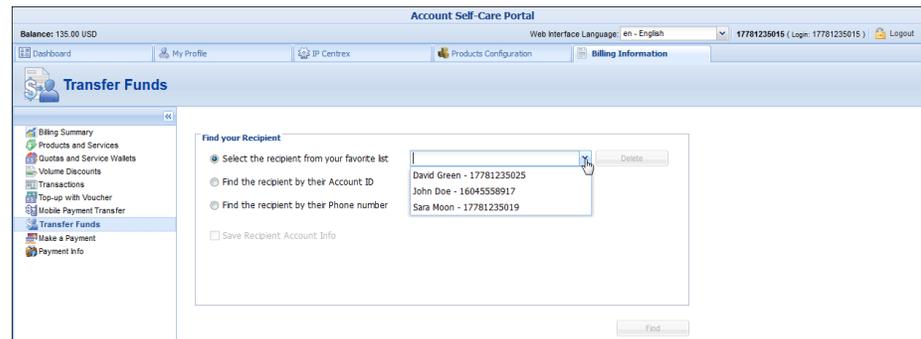
Service providers encourage end users to make frequent top-ups to provide for the transfer of airtime.

For instance, Jane Smith and her friend have identical service wallets that include Voice calls and Internet traffic. As the friend spent all his service wallet balance on Internet traffic, Jane Smith transfers 2 GB so they can continue chatting via the Internet. Afterwards, Jane finds out that her own service wallet ran out of money too, so she tops up her service wallet.

Favorites list

End users may create a favorites list that includes recipients they are accustomed to cooperating with. The list is easy to manage, and transferring funds and / or airtime can be done with just a few clicks. This saves time and makes the transfer process more user-friendly.

For example, returning to Richard Roe, let's say he visits the grocery store every week and pays for his goods by transferring funds. To have the grocery-store owner's contact information at hand, he includes it in his favorites list. Next time he goes grocery shopping he just selects the owner's account from the list and transfers the money. It's that simple and efficient.



Implementation specifics

In conclusion, the transfer of funds and airtime feature has the following implementation specifics:

1. Both sender and recipient must be allowed to transfer funds.
2. Only debit and credit accounts with individual credit limits can transfer and receive funds.
3. To allow accounts to transfer and receive funds, enable this option either at the customer or customer class level.
4. The currency of the sender and recipient must be the same.
5. Only accounts with identical service wallets can transfer and receive airtime.
6. To allow airtime exchanges, enable this option within the service wallet.

This solution enables service providers to increase their revenue by encouraging their end users to frequently top up their service wallets and their main balance to be ready for transferring money to their families and friends. Since no payment system is required to transfer funds and / or airtime, the service is widely used.

In addition to being an easy and secure way to pay for goods and services it encourages users to register new accounts. Thus, service providers also expand their customer base.

API to transfer airtime among service wallets and quotas

CSPs can allow their users to transfer airtime among their service wallets and quotas in such forms as GB, minutes, SMSs and money, etc.

Since the way that users transfer airtime may differ, the transfer is performed via the API. To do that, CSPs must develop custom self-care portals or mobile apps.

Consider the following specifics:

- The lifetime of transferred airtime synchronizes with the lifetime of the recipient's service wallet / quota.
- The transferred airtime is always used up first.
- To allow airtime transfer, enable this option within the service wallet / quota for both sender and recipient.
- The measured volume of service wallets / quotas must be identical (e.g. GB, minutes and money).
- Both sender and recipient must have the service wallet / quota assigned either to the product or directly to their accounts in PortaBilling®.

For instance, let's say that John Doe has 2 GB rolled-over quota from the previous month that expires on August 10. John opens his custom self-care interface and transfers these 2 GB to Mary's service wallet. Now Mary can use this quota until her service wallet expires, on August 28. Both John and Mary receive email notifications about the successful transfer.

To transfer airtime among service wallets and quotas using the API, use the `make_transaction` method of the Discount plan information service. For more information refer to the [PortaBilling® XML / JSON API Reference Guide](#).

With this feature, users can manage the airtime they've paid for and use it in full. This increases their loyalty and makes CSPs more competitive in the market.

Billing for time-based services

Call billing parameters

Traditionally, the billing of a specific user activity (e.g. a phone call) has been done only in the context of a particular session. This means that, when a call is completed, the billing engine calculates the charges for the call based on price parameters for the call destination and duration. This amount is then applied to the account, customer or vendor balance. This process does not take into consideration any other previous activity of the user (e.g. phone calls already made during this billing period).

Volume discount plans allow rating of service consumption based not just on information about a particular isolated event, but are able to

dynamically adjust billing with regard to previously consumed service (e.g. if you have made more than 500 minutes of calls to US&Canada this month already, the current call will be charged at a 10% discount to the standard rate).

Rating individual calls

Peak and off-peak prices

It is possible to have different sets of prices for peak and off-peak time. Off-peak periods are defined using the powerful and flexible Time::Period module. An Off-peak Period Wizard is also available, allowing you to perform period definition in the easy way. See the [PortaBilling® Help](#) for more details.

The fundamental concept here is the period definition, which specifies a certain interval in time (typically a repeating one). In each period you can specify conditions for:

- Time of day
- Day of week
- Day of month
- Month

If you do not specify a condition for a certain element (e.g. for **Month**) then it is considered that there are no limitations on this element, e.g. “from 8 p.m. to 8 a.m.” will apply in the same fashion throughout the year. If you define conditions for several elements, these must be satisfied at the same time, e.g. “from 8 p.m. to 8 a.m., Monday-Friday” means that an event fits into this interval if it happens between 8 p.m. and 8 a.m. **and** from Monday through Friday, e.g. 6 a.m. Saturday is excluded.

An off-peak period can be a combination of several definitions, so that a given moment is considered to be within the off-peak period if it satisfies *at least one* of these definitions. For instance, for the period “from 8 p.m. to 8 a.m., Monday–Friday OR Saturday–Sunday” both 9 a.m. Saturday (since it satisfies the second condition) and 6 a.m. Friday (since it satisfies the first condition) are included.

By default, everything is considered to be peak time. You may use one of the three available options to determine whether a peak or off-peak rate should be used. When defining an off-peak period, you may choose for off-peak rates to be applied if:

- a call starts during off-peak time.
- a call finishes during off-peak time.
- a call both starts and finishes during off-peak time (the least flexible option).

Multiple off-peak periods

It is also possible to offer two distinctive off-peak periods (e.g. different prices for calls made overnight or on weekends). You can define two independent off-peak periods by enabling the `MultiOffPeakPeriod` configuration option. In order to add a second off-peak period definition in the Off-peak Period Wizard, create your first off-peak period, then use the **Add another period** button. In order to use different rates for two off-peak periods, you must define prices in tariff rates for first and second off-peak periods.

So, for instance, in order to use the 0.10 rate during peak hours, 0.08 during weekends and 0.06 during night time, you should specify the main off-peak period as “8 p.m.–8 a.m.” and the alternative one as “Saturday and Sunday”, and enter the following in the rate data:

- Peak price_first, price_next = 0.10
- Off-peak price_first, price_next = 0.06
- Second off-peak price_next, price_next = 0.08

Charging calls – traditional method

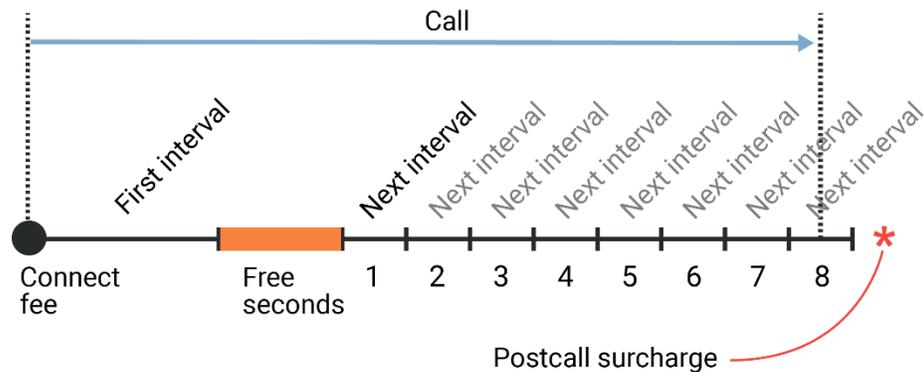


Figure 3-1

The figure above demonstrates how calls are charged. A **Connect Fee** is charged immediately upon connection, and all calls shorter than the *First interval* will be rounded to **First interval** seconds. Free seconds are granted after the *First interval*, so this part of the call is not charged. Calls longer than (*First interval*+*Free seconds*) will be rounded up to multiple **Next interval** seconds. After that, the **Post Call Surcharge** is applied. The call illustrated in the figure above will be charged using the following formula:

$$\begin{aligned} \text{Amount_Charged} &= \\ &= (\text{Connect_Fee} + \\ &\quad + \text{First_Interval} * \text{Price_First}/60 + \\ &\quad + 8 * \text{Next_Interval} * \text{Price_N}/60) * \\ &\quad * (1 + \text{Post_Call_Surcharge}/100) \end{aligned}$$

Parameters such as **First interval**, **Next interval**, **Price First** and **Price Next** can be specified per destination. **Connect Fee**, **Free Seconds** and

Post Call Surcharge are defined on a per-tariff basis, and so will be the same for all destinations within a tariff.



Attention: The billing unit may be any length, but price must always be entered on a per-minute basis. This allows better operations with tariffs, for example, comparing two tariffs with different billing intervals or entering rates into the system from an external source.

Charging calls – rating formula method

This method gives you maximum flexibility for call rating. You create a call formula and then apply it to the call. This formula can consist of an unlimited number of **charge elements** (the only limit is the length of the formula text). These elements are interval, fixed surcharge and relative surcharge.

- **Interval** – is the charged part of a call’s duration. The parameters for this element are:
 - **Count** – an increment that consists of one rounding period.
 - **Duration** – a rounding period measured in seconds.
 - **Price per minute** – the cost charged per call minute.

The charge for an interval is then calculated as follows:

Interval total = Count * (Duration/60) * Price per minute.

For example: You configure the interval to have the following parameters:

- Count – 3
- Duration – 10 seconds
- Price per minute – \$ 0.10

Thus, the interval total is \$0.50 as $3 * (10/60) * 0.10 = 0.50$.

Since a *count* is an increment consisting of one rounding period (i.e. 10 seconds), the number of counts is defined by the call duration. A count is rounded up when the call exceeds the interval duration. For instance, a call lasting 9 seconds is considered to have 1 count, while a call lasting 13 seconds has 2 counts.

Therefore, if a user makes a call that lasts 35 seconds, PortaBilling® calculates the charge according to the following parameters: duration – 10 seconds, counts – 4, price per minute – \$0.10. Therefore, the user is charged \$0.67 for this call.

Accordingly, you can create intervals such as “First 10 minutes – price per minute is \$0.05 and we will round it up to 30 seconds” –

so in this case, the price per minute will be 0.05, the duration 30, and the count 20 (because there are twenty 30-second intervals in 10 minutes).

- **Fixed Surcharge** – is a fixed amount added to the call cost at the moment when this amount is applied, namely, following a certain interval.
- **Relative Surcharge** – is a percentage of the value (the total sum of the charges for preceding intervals and previously applied surcharges) to be added to the call cost. This amount can be added either at the moment it is applied (following a number of charged elements) or at the end of the call.

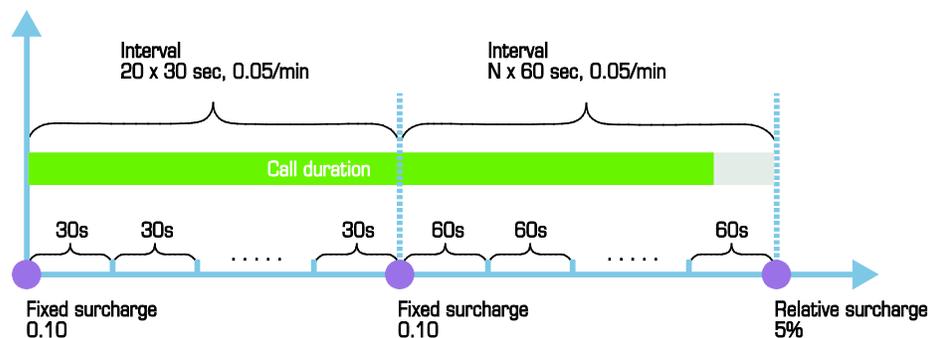
- **Call Disconnect** – is an option for forcibly disconnecting calls. Formula elements are applied until there is no more non-charged call duration left. The surcharges which follow a certain interval are applied only if this interval has been fulfilled. “Fulfilled” means that the interval is covered completely, i.e. the remaining call duration was enough to cover the $\text{Count} * \text{Duration}$ time. Let’s take an example; if we construct the formula:

- $3 * 60$ seconds, 0.10/min.
- Fixed surcharge 0.05.
- $N * 60$ seconds, 0.10/min.

then, when a call is made for 1 minute and 5 seconds, the charged amount will be 0.20 – in the first interval we use two increments of 60 seconds each, with price per minute 0.10. Since this first interval is to contain three increments, and we have used only two, the surcharge which follows the interval is not applied.

When a call is made for 4 minutes and 20 seconds, the total charged amount will be 0.55: three increments of 60 seconds each, a surcharge of 0.05, and then another two increments of 60 seconds with price per minute 0.10.

There is one exception to the surcharge application rule above: if the last formula element is a surcharge, it is always applied. This allows emulation of the “post-call surcharge” from earlier releases.



The figure above demonstrates how calls are charged. For example, the rating scheme “Charge 0.10 connection fee, then first 10 minutes should be rounded up to half a minute with price per minute 0.05. If the customer calls for more than 10 minutes, apply 0.10 fee and round the rest of the call up to 1 minute with price per minute 0.05; apply 5% post-call surcharge at the end of the call” will be translated into the following formula:

- Fixed surcharge 0.10.
- Interval 20 x 30 seconds, price per minute 0.05.
- Fixed surcharge 0.10.
- Interval N x 60 seconds, price per minute 0.05.
- Relative surcharge 5% .

Most often you would like to set up a general call rating scheme (rounding, surcharges, etc.) immediately, even though you foresee that your rates per minute will change in the future. In this case, you may construct the formula as you would normally do, but instead of entering a fixed price per minute in the formula you can specify: “use the current value of the **Price_Next** parameter”. Later, when you decide to change your rates, it will be much easier to download the current rate data, alter just one column (**Price_Next**), and upload the file again, rather than change the formula in every row. Also, using the pseudo-variables **First Price** and **Next Price** is required to ensure that the correct rate is used during the off-peak period, since the system will then automatically populate the variables with the values you specified for either peak or off-peak prices. Of course, if you entered a constant price value into the formula it will remain unchanged regardless of the off-peak period.

Using formula parameters linked to rate values has another advantage: if you have different prices or interval durations in peak/off-peak periods, the system will automatically use the one applicable.

Too short calls

Unfortunately this is a common problem: for some destinations (usually where analog lines are used) your vendor will report the call as connected once the remote phone starts ringing. So, even if no one picks up the

phone, such calls will still be considered successful by your system, with a duration (for instance) of 15 seconds, and the customer will be charged for them – which will probably lead to disputes and arguments. Using the new call rating formula, you can specify that short calls to a specific destination be disregarded. So, if you set the **Do not bill calls shorter than** parameter to 20 for a certain destination, and if a call lasts less than 20 seconds, it will not be billed at all. However, if the call was made for 20 or more seconds, it will be charged for the full call duration.

Payback (reverse) charging

A rate that is marked as **reverse** contains exactly the same set of parameters as the normal rate. The total amount is calculated according to the same rules as described above – with the exception that this amount is then **credited** to the customer’s account.

This allows for “awarding back” money to the customer for certain types of service (e.g. when he receives an incoming call).

The only difference between ordinary rating and rating according to reverse rates is that volume discounts do not apply for transactions produced according to reverse rates (the current discount has no effect on the credited amount, and volume discount counters are not updated for such transactions).

Billing by ANI (CLI)

While calls are normally rated based on the destination number, in some situations (e.g. charging the owner of a toll-free number for incoming calls) the caller’s number needs to be used to calculate billing charges (e.g. toll-free number costs are higher if the call originates from Alaska rather than from the 48 contiguous states).

Instead of using the default mode, where the billing engine looks up a matching rate based on the destination number (also called CLD or DNIS), it is possible to use an alternative method where the caller number (also called the CLI or ANI) is used to find the rate.

In calculating CDRs for vendors, the rating switch is done in the connection settings. So you can have some “normal” connections to a vendor, where calls are charged based on CLD/DNIS, and others where the rating is done based on the caller number (CLI/ANI). Also part of the connection configuration is the access code, which is used to apply a tariff to the customer who receives a call (Please see more details in the **Unified PortaSwitch Handbook Collection** about the specific functionality available.).

For customer billing the configuration is done in the product. A VoIP from vendor connection which will initiate billing by ANI/CLI contains

the access code used to select the applicable tariff in the product when the customer is billed. Then, in the **Usage Charges** configuration of the product, the administrator will add an entry with that access code and set the **Rate Match Mode** parameter to “Calling Number”. This allows the use of one product that will contain both “normal” tariffs (applied to outgoing calls when the call is billed by CLD) and “ANI-based” tariffs (applied to incoming DID calls).

Prepaid card billing features

PortaBilling® offers an extremely flexible rating and authorization engine for providing prepaid calling card service. It contains the most comprehensive set of special features required for this type of service. Please see more details in the [Unified PortaSwitch Handbook Collection](#) about the specific functionality available.

Session billing parameters

Time-based services such as WiFi access, where the cost depends on the total session duration, are rated exactly the same as voice calls.

Always-on services

Always-on services are those which are consumed without interruption over extended periods of time. The only thing that varies is the rate of consumption. For instance, a computer connected to the Internet is constantly transferring data. However, the amount tends to be very low at night, when no one is using the computer, and higher during the day, when emailing or web browsing are taking place. It can even be extremely high when content like video-on-demand is being downloaded.

The biggest challenge when charging for services such as data downloads is that, although the customer uses the service all the time, some specific, “finite” entries must appear on his invoice, e.g. “Data downloaded May 1st–May 31st.” Also, customers expect to see real-time information about the charges applied, as in the case of other types of services.

Normally, PortaBilling receives information about service usage via “keep-alive” (or “update”) requests. Each of these informs billing that the session is still in progress and what the current service consumption is (e.g. the amount of currently downloaded and uploaded data). For each such request, PortaBilling determines which billing session it relates to. For example, although from the customer’s point of view the session did not terminate at midnight on June 1st, this was nonetheless the end of his billing period, and a new billing session has started, with charges accumulating to a new xDR. When processing an update request, PortaBilling calculates the new amount to be charged for session activity up to the present moment, and updates the xDR for the session accordingly. This billing method, i.e. constantly modifying an xDR record

to reflect the new charged amount, is quite different from the one used for other types of services, where PortaBilling receives the information only after the service has been completely consumed, and so the xDR is created only once.

PortaBilling's open architecture allows to process virtually any type of "always-on" service. Take electricity billing, for example: your household consumes electricity non-stop, but consumption is lower during the night (when only the refrigerator is running), higher during the evening (when the lights are on), and extremely high during the day (when various appliances are in use).

Flexible definition of the party to charge

Modern business involves diverse types of services, which require different types of billing. Calls made by prepaid cards are billed by PIN number, calls from a customer's GW should be billed by IP address, some wholesale customers are identified only by the port they are connected to, and so on.

It is easy to find out that a billing event occurred (e.g. a call was made). The important thing is to correctly determine who made the call and, therefore, who should be charged for it. PortaBilling® uses a simple and flexible yet powerful method to determine a call's owner. Identification of the account (Account ID) is sent to the billing system in the **User-Name** attribute.

For example, in the case of prepaid calling cards, **User-Name** will contain the PIN, IP address for IP-based billing, and so on. Of course, an account with such an account ID must exist in the database. This way of utilizing the **User-Name** attribute is actually default behavior for Cisco call handling applications, and the latest firmware releases of Quintum should do the same.

Since PortaBilling® supports account aliases, in some cases the value of the **User-Name** attribute will first be matched against the ID of an alias, which then will be used to retrieve the main account information. However, the principle remains the same.

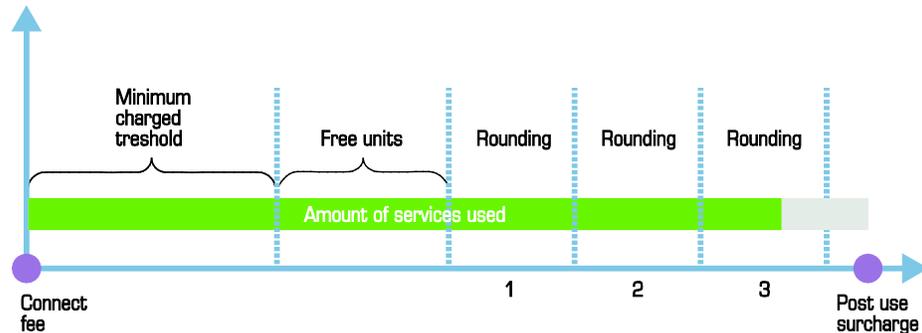


Due to such an approach, logic in the billing is independent from service-specific issues. You are not limited to any pre-defined billing schemes in PortaBilling®; rather, you can design your own. Today you can bill by PIN or IP address, and tomorrow you will additionally be able to bill by `h323-id`. All you need to do is pick the proper application to handle the call, so that a correct identification will be sent in the **User-Name**.

Unfortunately, for some types of service User-Name is not sent to the billing properly on Quintum gateways. In this case, it is possible to employ username replacement rules on the billing side to emulate this. See the relevant section in the "How to:" chapter for more details.

Billing for quantity-based services

Quantity-based services (such as the amount of data transferred, SMS or IPTV purchases) are billed rather similarly to the “default” rating of calls, with various rating elements applied to the “total used service” value. In fact, if we regard time-based services as those where a quantity of seconds is used as the rating parameter, then the two are almost identical!



The picture above demonstrates how quantity-based services are charged. A **Connect Fee** is applied for all such billing events. For any event where a customer has used less than the *Minimum Threshold* units, he will be charged as if he used the *Minimum Threshold*. *Free Units* are granted after the *Minimum Threshold*, meaning that the number of units specified in this parameter will not be charged. The amount remaining after (*Minimum Threshold* + *Free Units*) will be rounded up to multiple **Rounding** units. Note that at this point only measurement units are being used, and these must be converted to billing units, so that the actual amounts charged will be calculated as $Price_per_unit * Units / Billing_Ratio$.

After the total charged amount has been computed, the **Post Use Surcharge** is applied. The service illustrated in the figure above will be charged using the following formula:

$$\begin{aligned} \text{Amount_Charged} &= \\ &= (\text{Connect_Fee} + \\ &\quad + \text{Minimum_Threshold} * \text{Unit_Price_Initial} / \text{Billing_Ratio} + \\ &\quad + 3 * \text{Rounding} * \text{Unit_Price_Next} / \text{Billing_Ratio}) * \\ &\quad * (1 + \text{Post_Use_Surcharge} / 100) \end{aligned}$$

Parameters such as **Minimum_Threshold**, **Rounding**, **Unit Price Initial** and **Unit Price Next** can be specified per destination. **Connect Fee**, **Free Units** and **Post Call Surcharge** are defined on a per-tariff basis, and so are the same for all destinations within a tariff.

To understand this better, let's take another simple example. A data transfer service rating is defined as follows:

- The measurement unit is a byte, the billing unit is a kilobyte, and the billing ratio (how many measurement units make up one billing unit) is 1,024.
- The rate is defined as follows: minimum threshold – 10,240, rounding – 1,024, and the initial and next prices are the same – \$0.02. (Prices are defined per *billing* unit, so this is a price per *kilobyte*.)

So if a customer initiates a session and transfers 1,976 bytes, this will fall under the minimum threshold; thus he will be charged $10,240 * 0.02 / 1,024 = \0.20 . If the customer transfers 17,290 bytes during the session, this results in 7 rounding increments above the minimum threshold, and so the total charged amount is $10,240 * 0.02 / 1,024 + 7 * 1,024 * 0.02 / 1,024 = \0.34 .

Overdraft protection

The traditional authorization scheme, in which the system checks the amount of available funds at the beginning of a call, has many drawbacks. For one, the system does not respond to balance changes (refunds, payments, or charges) made while the call is in progress. Since every session is authorized based on the total amount of resources available, overdrafts are possible in the case of multiple concurrent sessions and the like. The alternative – restricting the number of concurrent sessions to one – allows you to avoid overdrafts, but is inconvenient for the end user.

This is why PortaBilling® offers you the ability to dynamically lock the funds required to pay for a certain amount of a service before that service is used. Proper implementation of overdraft protection requires support both on the billing engine side (PortaBilling®) and on the part of the network element providing the service, e.g. PortaSIP®. (See below for a description of the scenarios applicable to “legacy” switches.)

Authorization and fund locking

What is the relationship between authorization and fund locking? Authorization is simply a check for whether currently available funds are sufficient to cover a service (or how much of a service can be used with the funds currently available). Using a simple analogy, it is like taking a look in your wallet to see how much money you have left before deciding whether you have enough to purchase a DVD or buy some chocolate bars. The problem is that such a check is done without taking into account your other attempted purchases. For example, if you have \$20 available, the system will tell you that you have enough money to buy 1

DVD (\$15) or 10 chocolate bars (\$2 each), but it cannot tell you whether you can buy a DVD and **also** buy 2 chocolate bars.

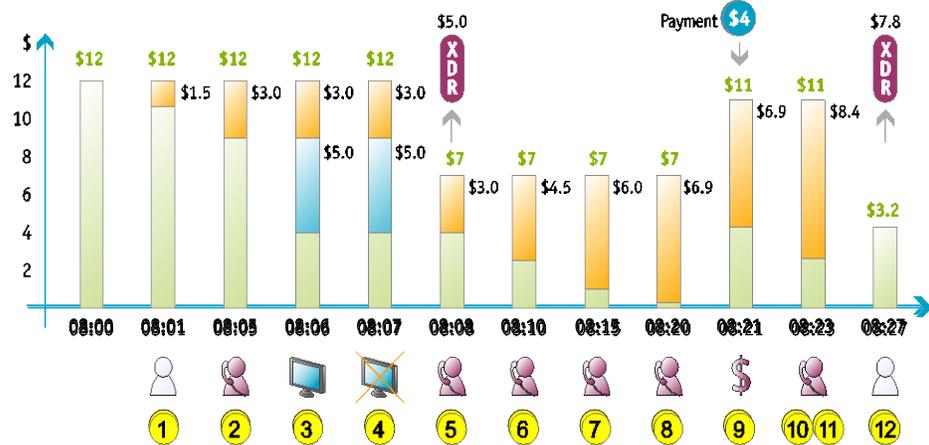
Fund locking not only checks that you have enough money for the DVD, but also immediately puts that money into a different compartment of your wallet. Now when you check to see how many chocolate bars you can buy, you will be working with only the remaining portion of the money. This will spare you an awkward situation at the cash register.

Dynamic authorization

The simplest overdraft protection scenario is “resource locking.” This is done when an application knows the amount of funds required beforehand. So right at the start, the application sends a request to validate the availability of these funds and if this is successful, it locks them. Then, when the purchase has been completed, a “charge” event is sent to PortaBilling. One example of such a service is pay-per-view content distribution. After a customer clicks on a “Get the movie” link, the system authorizes and locks the funds to cover the cost of the movie to ensure that the customer has the right permissions and sufficient funds. The customer is then shown a confirmation dialog box. If he approves the purchase, the money is deducted from his account

Dynamic re-authorization is a more complex process used when it is not possible to tell how much a customer will spend in a particular session, so that additional funds must be locked “on the go.” When a session (e.g. a voice call) starts or is in progress, the system “reserves” a small portion of available funds to cover the next time interval (e.g. 5 minutes). As the end of this interval approaches, the system attempts to reserve funds to cover the next interval, etc. If no remaining funds are available, the call is disconnected. So if a customer sends an SMS message during a phone call (decreasing the available funds), or makes a credit card payment (increasing the available funds), that immediately affects the call – and either the call is disconnected sooner or the maximum allowed call time increases.

Let’s see how dynamic re-authorization works in a specific example where an account is concurrently using multiple services and making a payment while the session is still in progress.



- Initially, the account has \$12 available.
- The user starts a phone call to a destination with a price of \$0.30 per minute. PortaSIP® requests call authorization for 5 minutes. At this point the user’s balance is still unchanged, but funds covering a 5-minute phone call (\$1.50) are locked (1).
- While the call is in progress, at some time before the end of the “approved” period (for our example, we’ll assume this is one minute), PortaSIP® will attempt to perform re-authorization.
- PortaSIP® sends a re-authorization request to billing. PortaBilling verifies that the account has enough funds for another 5 minutes. The amount locked is now \$3.00 (2).
- While still on the phone, the user decides to purchase a pay-per-view movie. An authorization request arrives from the IPTV platform, and the funds needed to pay for the movie (\$5.00) are locked in addition to the previous \$3.00 (3).
- The user now attempts to purchase another movie (\$5.00), but this request is rejected, since he only has \$4.00 left (4).
- The purchase of the first movie is completed, and the balance decreases to \$7.00 (5).
- Since the call is still in progress, and the end of the “authorized” time is approaching, PortaSIP® sends another authorization request, and so funds for another 5 minutes are locked (a total of \$4.50 is now locked) (6).
- The user keeps talking, and so 5 minutes later yet another authorization is sent. The total amount of locked funds is now \$6.00 (7).
- Soon another authorization is sent, but now there are only enough funds for 3 minutes. \$6.90 is now locked (8).
- The user realizes that he is running out of money, but wants to continue the phone call, and so he promptly makes an online payment with his credit card. \$4.00 is deposited to his account (9).
- Since the last re-authorization provided less time than requested (3 minutes instead of 5), PortaSIP® determines that the user is out

of money, and normally it would disconnect the call. However, it makes one last re-authorization attempt, one minute before the call is to be disconnected (10).

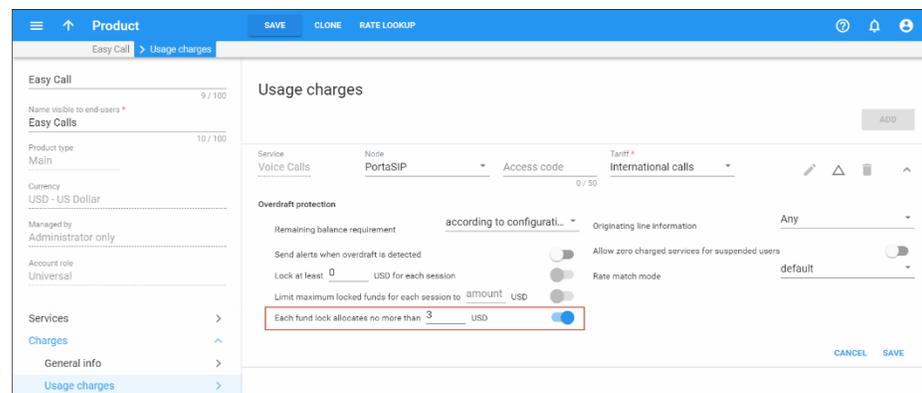
- Since the user now has available funds, another 5 minutes of call time are authorized. The total amount of locked funds is now \$8.40 (11).
- Finally, the user disconnects, and is charged \$7.80 for the full call duration of 26 minutes. Thus his balance is now \$3.20 (12).

Limiting the maximum fund amount locking

There are two situations that may lead to an excessive locking of funds:

- Attempting to use an “expensive” service – For instance, let’s assume that PortaSIP® is configured to lock sufficient funds for a 15-minute conversation at the start of a call. This is fine for calls to the US or Europe (in which case only 20-40 cents would be locked). But if a customer with a fairly low balance calls an expensive destination (e.g. Somalia – \$1.00/min), he may find that all of his funds are suddenly locked and he cannot use another service in the meantime.
- Old-style applications that do not support re-authorization – If an authorization request comes from a gateway or switch that does not support dynamic fund locking, then all available funds will be locked while that session is in progress.

To prevent locking too many or all available funds, the administrator can limit the maximum fund allocation for each subsequent fund lock.



The screenshot shows the 'Usage charges' configuration page in the Porta Billing administration interface. The page is for 'Easy Call' and shows various settings for 'Usage charges'. A red box highlights the 'Limit maximum locked funds for each session to' field, which is set to '3' USD. The 'Each fund lock allocates no more than' field is also highlighted with a red box.

For “overdraft protection-aware” applications, each request will lock no more than the amount specified in the **Each subsequent fund lock allocates no more than** field. Taking our example of a call to Somalia, if we set this parameter to \$3.00, the billing engine will respond to PortaSIP’s request to lock funds sufficient for 15 minutes by locking up only \$3.00 (3 minutes). Thus PortaSIP® will perform a re-authorization every three minutes.

For applications that do not support dynamic re-authorization, the maximum call duration will be limited to the amount defined in the **Each subsequent fund lock allocates no more than** field. Thus, if we set this parameter to \$3 and dynamic re-authorization is disabled for the call to Somalia, it will last only 3 minutes and then be disconnected.

Volume discount plan counters locking

By using volume discount plans, service providers can either adjust users' pricing depending on the volume of the service used or they can allocate some free services (e.g. minutes, GB, etc.) to their users. In cases where several accounts share the same quota and initiate sessions more or less concurrently, PortaBilling® must be aware of quota usage by each account to prevent them from going into overdraft.

Therefore, when authorizing a session, PortaBilling® allocates a portion of the available quota and locks that allocated portion for the authorized session. Then it excludes that locked quota portion when it authorizes simultaneous customer sessions.

When locking a quota, PortaBilling® considers the following:

- customer balance;
- tariff rates;
- maximum session duration either defined by the administrator or provided by NAS; and
- the fund locking configuration defined within the product.

Consider the following example:

Let's say John Doe has a \$0 balance and a \$10 credit limit. His tariff plan allows him and his wife Jane to make calls to Canada for \$0.5 per minute. They also share a quota for 50 free minutes to Canada. For them to make simultaneous calls and at the same time, not exceed their credit limit, PortaBilling® authorizes their calls for \$3. The administrator has defined that all calls may last no longer than 30 minutes.

When John calls 16045556754, PortaBilling® considers the available quota against the maximum allowed call duration – 30 minutes. Since $50 > 30$, PortaBilling® locks 30 minutes from the quota, and leaves 20 minutes available.

Five minutes later, Jane dials 16045557785. PortaBilling® considers the remaining quota – 20 minutes, against the maximum allowed call duration – 30 minutes. Since $20 < 30$, PortaBilling® locks the remaining quota and adds another 6 minutes to it, calculated based on the call cost and the authorization amount ($\$3 * \$0.5/\text{min} = 6 \text{ min}$). Thus, Jane is authorized to make a 26-minute call.

After 15 minutes, John hangs up. Since John's call lasted 15 minutes, PortaBilling® unlocks those unused 15 minutes and makes them available for other calls.

PortaBilling® locks the counters for any type of volume discount plan (discounts, quotas and service wallets) for both RADIUS- and Diameter-based sessions. To enable this functionality, your service configuration must meet the following conditions:

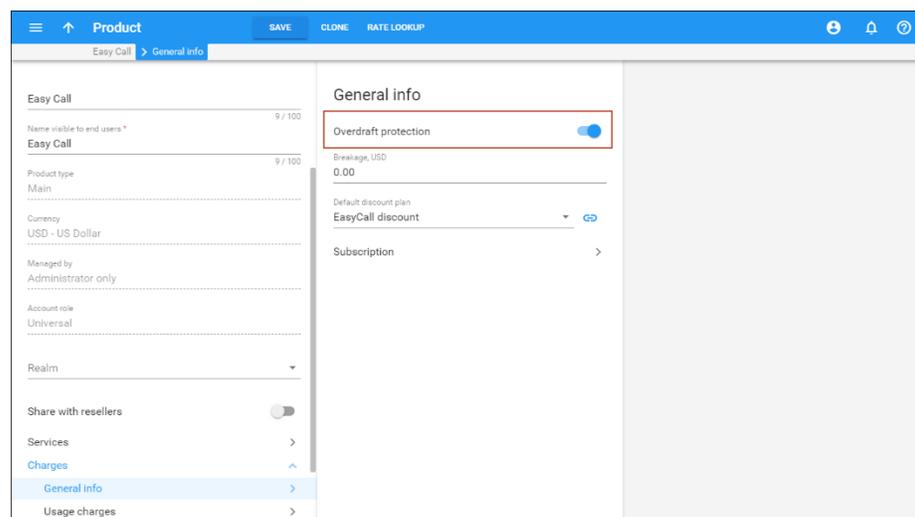
- Your product must have **Overdraft protection** functionality enabled; and
- The value of the discount threshold must have a limit.

Overdraft protection with legacy equipment

In order to fully use overdraft protection capabilities, the node providing the service (gateway, switch, etc.) must send the proper attributes in the RADIUS request to PortaBilling®. During the initial authorization request, it will inform PortaBilling® of the desired interval for fund locking (and the minimum possible interval for re-authorization), and will then periodically issue re-authorization requests. Real-time overdraft protection is fully supported by PortaSIP®.

What happens when a node that is not capable of dynamic re-authorization sends a “standard” RADIUS authorization request? Also, how does fund locking affect the number of concurrent sessions available for a customer?

Answer: the product configuration features an **Overdraft Protection** setting that allows you to define how fund locking is applied:



- When overdraft protection is **disabled**, that means there are effectively no “locked” funds. Thus, for concurrent sessions, PortaBilling® computes the maximum session duration based on the total amount of customer’s available funds for each session separately. So in the case of several concurrent sessions, each can potentially use up the available funds. This method provides transparency for the application using overdraft protection, although in reality, it does not provide greater security than a simple authorization request does. For obvious reasons, it is not recommended for general use.

NOTE: When disabling overdraft protection functionality, make sure that the product is only assigned to credit accounts. Otherwise, it may open the overdraft option for debit accounts.

- When **enabled**, fund locking is done on all account types. For “overdraft protection-aware” applications, the amount of funds locked is the lesser of either the amount necessary to cover the requested service or the amount defined in the **Each subsequent fund lock allocates no more than** field. For legacy authorization requests, it represents either all of the available funds or the amount defined in the **Each subsequent fund lock allocates no more than** field. Note that if all of an account’s available funds are locked for a session, more paid sessions cannot be established.

4. Recurring charging

Subscriptions

The **Subscriptions** module allows you to charge periodic fees to your customers for using the service. If your advertisement states something like “only \$9.99 per month, and just \$0.99 extra per month for Voicemail service,” then two subscription plans are involved. The parameters you specify within a subscription plan affect the customer balance. Each subscription plan includes the following:

- Name and description – Your administrators use this to better manage various subscription plans.
- Subscription plan name visible to end user – This is a clear subscription name that is shown to end users on invoices and on their self-care interfaces.
- Activation mode – This specifies when a subscription is active and when charges start. A subscription is typically activated upon its creation (or on the start date, if that is specified). PortaBilling® supports an additional mode of subscription activation based on the account’s first use. You can avoid issues when there is a delay between the time a customer signs up for a service and the time they can actually use it (e.g. the customer uses online signup to purchase a residential SIP service, but their IP phone is not delivered by FedEx until five days later).
- Minimum subscription period and early cancellation penalty – It is common practice to lock a customer’s contract for a certain period (e.g. if you provide the customer with a free IP phone, you want to make sure they continue paying monthly fees until the cost of the phone is paid back). If the subscription service is canceled earlier than the interval specified, the customer is charged a cancellation penalty.
- Subscription fees – These are the charges applied while the customer has an active subscription; see below for more details.

Subscription fees

When a customer’s billing period closes, PortaBilling® calculates the charges for any subscriptions that were active during this period. For each subscription, there are three types of charges:

- Activation fee
- Cancellation fee
- Periodic fee

Activation fee

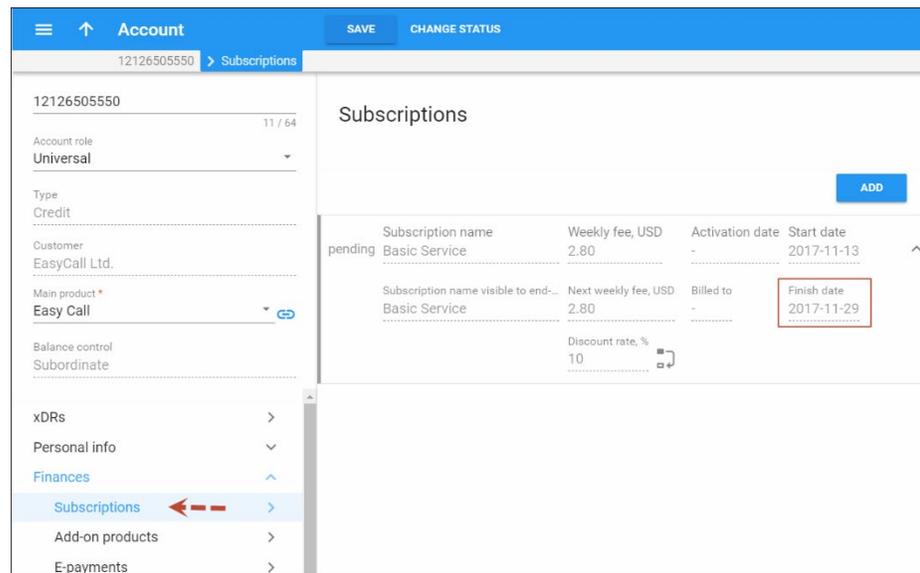
An activation fee automatically applies when a subscription start time begins, with one exception. Namely, in the subscription plan’s properties

you may specify that the subscription plan will be active only when the account is used for the first time. In this case, the subscription is activated either on its start date or on the date of the account’s first use, whichever comes later.

Cancellation fee

A subscription which is not yet active can simply be deleted (for instance, a customer signs up for a new service beginning next month, but then changes his mind). A subscription that is already active cannot be deleted (see the [Assigning subscriptions](#) section for details). When a subscription is canceled, the cancellation date is entered as the finish date for the subscription and the customer is charged a cancellation penalty.

The cancellation date is typically the day on which cancellation takes place. However, you can also schedule cancellation for a future date (e.g. a customer wishes to use the service until the end of the present month)..



The record of the canceled subscription allows PortaBilling® to correctly charge for closed subscriptions. For example, if a customer subscribes to the service on the 3rd and cancels it on the 7th, he is still charged for five days of service.

You can define the cancellation penalty as either a fixed amount or the remaining subscription charge. In the latter case, if the customer’s \$5 subscription must be active for 10 months but they cancel it in six months, the customer’s cancellation penalty is \$20, which is the sum of charges for the remaining four months.

Note that if you change the subscription fee (e.g. from \$5 to \$7) before a customer cancels their subscription, PortaBilling® uses the new fee to

charge the cancellation penalty. Thus, a customer from the example above is charged \$28 ($\$7 * 4$ months) as a cancellation penalty.

Periodic fee

A periodic fee defines the recurring charges applied to a customer for service usage.

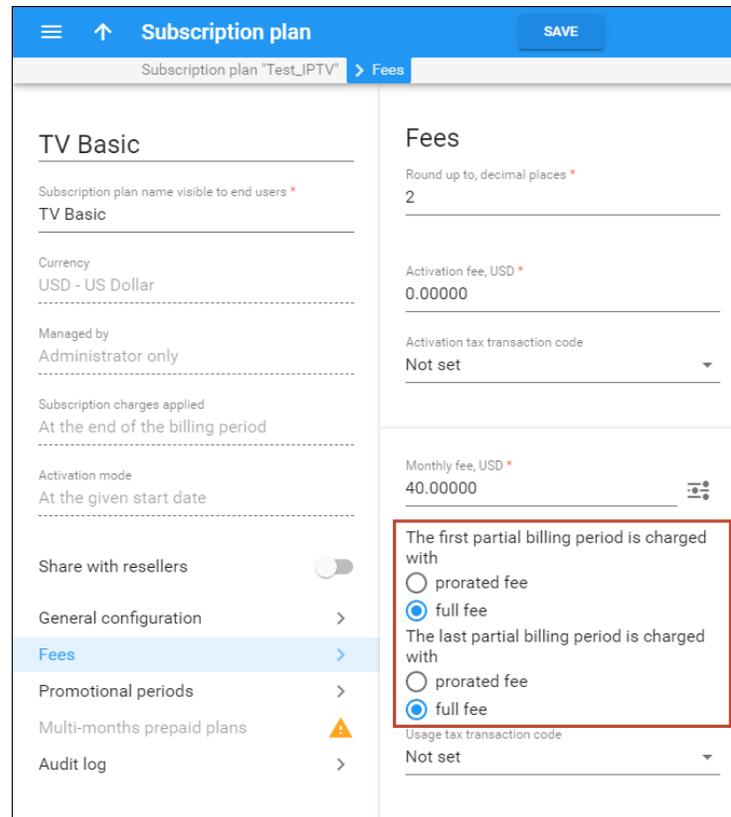
When a customer signs up for or cancels a subscription in the middle of a billing period, you can charge them either the prorated or full periodic fee. Thus, if a customer with a monthly billing period signs up for the subscription on March 17th, the charge can be either for 15 days in March (17th-31st) when the subscription is active or for the full month.

The following examples relate to the same subscription plan with a monthly periodic fee of \$9.99, and illustrate how it is prorated:

- Customer A, with a monthly billing period, activates the subscription on April 12th. On May 1st, the customer is charged \$6.33 because they used the service for 19 days (April 12-30th) and $19 * (\$9.99/30) = \6.33 .
- Customer B, also with a monthly billing period, activates the subscription on April 12th but cancels it on the 25th. On May 1st, the customer is charged \$4.66 because they only used the service for 14 days (April 12- 25th) and $14 * (\$9.99/30) = \4.66 .

These types of subscription fees are reflected on the invoice as extra invoice lines.

By default, first and last partial billing periods are charged with a prorated fee. To charge for subscriptions in full, go to **Subscription plans > Fees** and select the **Full fee** option.



The screenshot shows the 'Subscription plan' configuration page for 'Test_IPTV'. The 'Fees' section is active, showing the following settings:

- Subscription plan name visible to end users: TV Basic
- Currency: USD - US Dollar
- Managed by: Administrator only
- Subscription charges applied: At the end of the billing period
- Activation mode: At the given start date
- Share with resellers:
- General configuration: >
- Fees: > (highlighted)
- Promotional periods: >
- Multi-months prepaid plans: ⚠
- Audit log: >

The 'Fees' configuration includes:

- Round up to, decimal places: 2
- Activation fee, USD: 0.00000
- Activation tax transaction code: Not set
- Monthly fee, USD: 40.00000

A red box highlights the following options:

- The first partial billing period is charged with:
 - prorated fee
 - full fee
- The last partial billing period is charged with:
 - prorated fee
 - full fee

Usage tax transaction code: Not set

Note that this option is available if you charge for subscriptions **At the end of the billing period** or **In advance** and is not available for **Progressive** (daily) subscriptions.

Also, if you already use the **Waive the charges for the days when no service was provided** functionality, note that there are no refunds for the first and last billing periods if you set up the non-prorated (full) fee.

The **Full fee** option enables you to adjust subscriptions to your business requirements. For instance, you may want to balance the cost and revenue while you resell your partner's services on a recurring fee basis.

Let's consider an example company ABC, which offers IPTV services to end users for a monthly subscription fee. The IPTV vendor charges ABC in full even if their customer cancels the subscription. To avoid revenue leakage, the ABC company can also charge the customer in full for the last month of service usage.

Thus, when charging full fee for partial billing periods, you can ensure that the subscription charges cover the cost, no matter when the subscription is activated / canceled.

Also, you can modify the periodic fee for a subscription plan once it has been activated (e.g. you can decide to increase / decrease the price for a service).

In this case, fee changes have no influence on the current billing period and are applied to the next billing period. In other words, there is no recalculation. The next billing period's charges use the newly modified periodic fees.

For example, if a subscription fee of \$10 is charged one month in advance and the current monthly billing period ends on April 10th, the customer is charged until May 10th.

On April 25th the subscription fee changes and becomes \$8. The customer's invoice (produced on April 11th) contains charges for all calls made prior to April 10th and subscription charges covering the period from April 11th to May 10th without any recalculation (i.e. the subscription charges are still \$10 for this billing period despite any changes).

The new subscription fee (\$8) is applied to the new billing period, which covers the period from May 11th to June 10th, and the customer continues being charged \$8 for subsequent billing periods.

Periodic fees configuration

Since your customers may have different billing periods, PortaBilling® provides you with the flexibility to individually assign subscription costs for each period. For instance, if your monthly rate is \$19.99, you may want to charge a higher weekly rate than \$4.99, since the maintenance required by customers with shorter billing periods increases, and justifies a higher rate.

For instance, PortaBilling® allows you to define your monthly rate as \$19.99, your semimonthly rate as \$10.99, your weekly rate as \$6.99 and daily rate as \$1.99 for the exact same subscription plan. PortaBilling® will automatically use the correct base value according to the customer's billing period.

Promotional periods

Yet another common business practice is to offer special rates for an initial period following signup (e.g. "Only \$9.99/month!") while the disclaimer states: "For the first six months only, after which the standard rate of \$29.99 / month applies." PortaBilling® allows you to define an unlimited number of promotional periods, with different subscription fees for each. For example, you could create a subscription that offers free service for the first three months, a rate of \$9.99 for the next nine months, and \$12.99 thereafter; or something even more complex.

Incompletely used promotional periods

A subscription's promotional period includes each billing period during which a service is used, regardless of how many days the service is used in a particular billing period.

For example, let's say you offer a promotional period for three months and your customer John Doe subscribes to it on July 15th. John Doe is charged for the second half of July (16 days) and for the following two months (August and September) according to the promotional rate. Starting from the fourth month (October), the default rate is used.

In order to avoid potential misunderstandings when a customer expects a longer promotional period (till October 15th from the example above), the best practice is to use an anniversary billing cycle, as this eliminates the problem of "incompletely used" promotional periods.

Applying subscriptions

Since a product defines the way you provide a service to an end user, and subscriptions define charges for this service, it is obvious that they must be interconnected. That's why subscriptions are usually included with products (they can be selected from the subscription list.) To define a flat rate for provided service(s) and avoid discrepancies between subscription plans with different configuration parameters, only one subscription plan per product (both main and add-on ones) is allowed.

Subscriptions that an account receives with the products requested are called obligatory subscriptions, because it is not possible to remove or cancel these subscriptions as long as the account uses these products. You may, however, add optional subscriptions to an account (e.g. if an EasyCall product user wishes to receive Voicemail service). These optional subscriptions can be added or canceled as you wish. Note that you can only apply the same subscription once to an account if it has not been explicitly defined in its configuration (see the [Multiple subscriptions](#) section for details).

You can also assign subscriptions to a specific customer. Here the procedure is exactly the same as with an account. Since a customer cannot have a product assigned directly to him, there are no obligatory subscriptions in this case, i.e. all of a customer's subscriptions are optional.

Add-on product subscriptions

Both a feature and a subscription are attached to add-on products, so an administrator must only assign an add-on product to an account. Once the feature is activated, the customer is automatically charged for it.

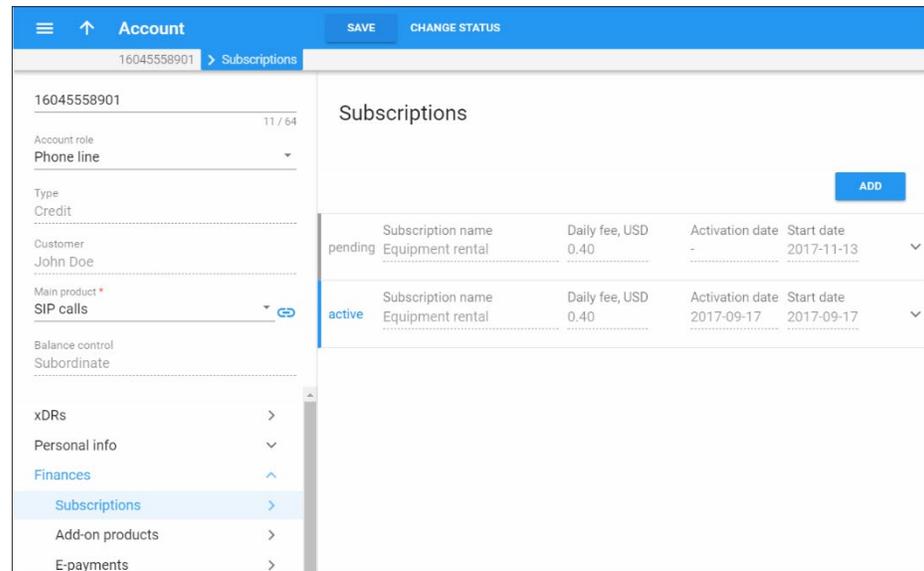
For example, a customer is assigned the “Call Waiting” add-on product with the call waiting feature enabled and a subscription associated with it that applies recurring charges for its usage. If the customer no longer needs call waiting functionality, the administrator removes the add-on product from the account properties with a single mouse click.

Multiple subscriptions

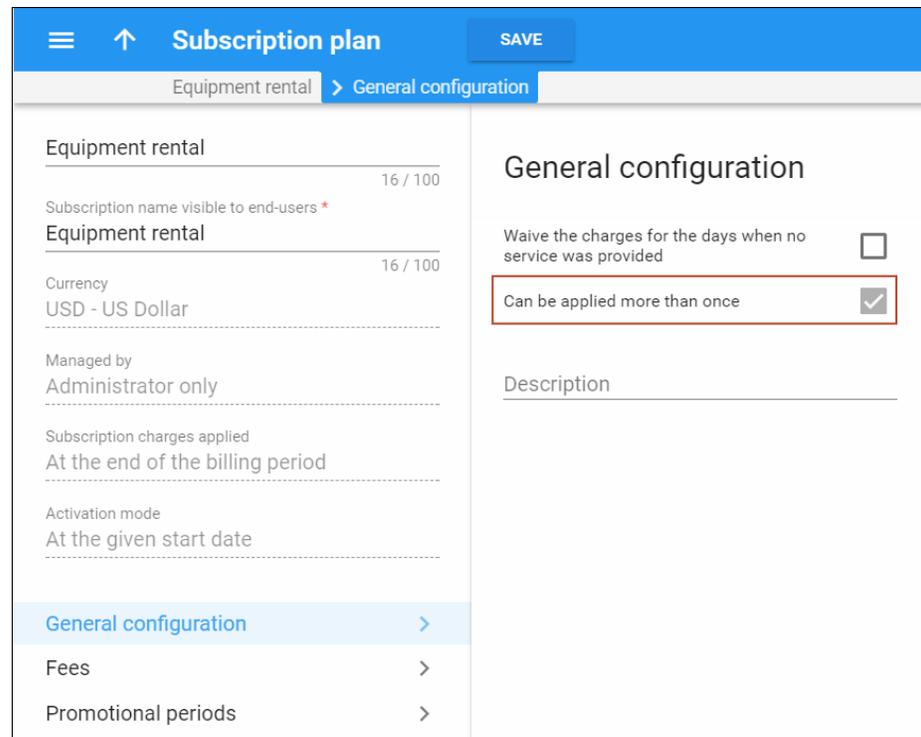
In most cases, every single subscription is charged only once even if this subscription is included in with several add-on products assigned to the account.

But in some cases it is necessary to assign the same subscription multiple times.

Let’s have a look at the following example. One of your customers is assigned a subscription for renting a router from you. On another day, they want to rent a second router and must therefore be charged the same subscription twice: for the (first) already rented router and for the newly rented one. Since this customer’s account already has a subscription for renting a router, a new add-on product is created with the same subscription assigned for the rent of the second router.



To make this happen, the administrator enables the “Apply more than once” option within the subscription plan configuration. This makes it possible to assign the same subscription to one account multiple times.



Subscription plan

Equipment rental > General configuration

Equipment rental

Subscription name visible to end-users *
Equipment rental

Currency
USD - US Dollar

Managed by
Administrator only

Subscription charges applied
At the end of the billing period

Activation mode
At the given start date

General configuration >

Fees >

Promotional periods >

General configuration

Waive the charges for the days when no service was provided

Can be applied more than once

Description

This model of applying subscriptions makes it possible to:

- provide revenue assurance, because using an add-on product guarantees that the feature will be activated, and the customer automatically charged for it; and
- save time, because everything can be done in one place using an add-on product.

Changing the subscription plan for a product

You may change the subscription plan for a given product at any time. This, however, affects all existing accounts that have this product. Once you select a new subscription plan for a product, the change takes immediate effect: the old subscription plan is removed and the new one is applied to *all* accounts with this product assigned..

Changing an account's product and subscriptions

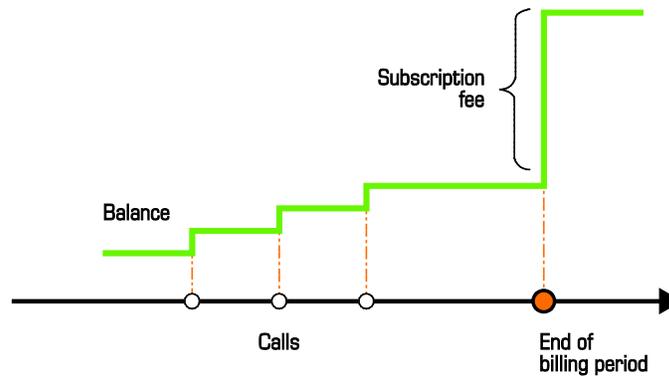
A change of product for an existing account, e.g. from product A to B, will have the following effect on this account's subscriptions:

- The obligatory subscription plan (inherited by the account from product A) changes to the subscription plan assigned to product B.
- Optional subscriptions (assigned to the account directly) remain unchanged.

Ways to charge for subscriptions

Traditional method

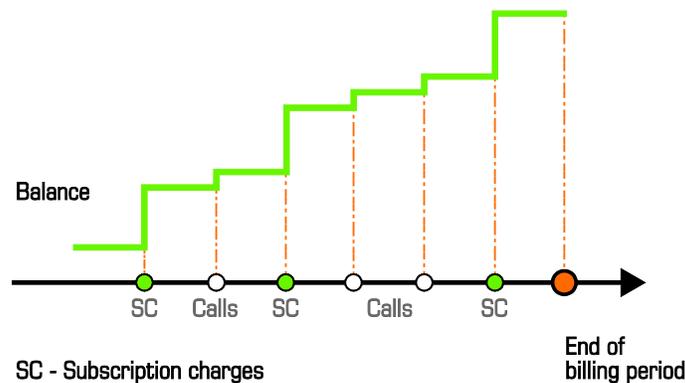
This is a very common and widely popular method. When a customer’s billing period closes, PortaBilling® calculates the applicable charges for all subscriptions that were active during that period and applies them to the customer’s balance.



Consequently, once the billing period is closed there may be a significant increase in the customer’s balance.

Progressive method

While the situation described above is acceptable for postpaid customers with a large credit limit, or if a customer has a credit card that is automatically charged when a billing period is closed, it may create certain problems for prepaid customers. If such a customer does not have sufficient available funds at the end of the billing period, his balance will exceed the credit limit and his outgoing calls may be blocked. PortaBilling® provides a solution to this problem by allowing you to conveniently offer prepaid VoIP services in combination with subscription fees.



To avoid the unpleasant situation of a sizeable charge at the end of the billing period, PortaBilling®'s progressive charges are continuously applied. This means that at any given moment in time, the charged amount will cover the interval from the beginning of the billing period until the current day. For example, a customer with a monthly billing period signs up for a subscription with a \$9.99 monthly fee on April 1st. This means that every day of the billing period the customer is charged $\$9.99 / 30 = \0.33 . By the end of the billing period the total charged amount is equal to \$9.99. In effect, the customer is charged small portions of the full fee every day. His balance grows slowly, giving him enough time to act in response and deposit more money into his account.

Subscriptions charged in advance

This mode of charging allows service providers to charge customers for one or more billing periods in the future. PortaBilling® applies 'in advance' charges for the following period at the end of the current billing period.

Let's say that your postpaid customer John Doe wants to rent an IP phone from you. You want John to pay for the phone rental in advance so you assign a subscription plan that is \$30 with a \$10 activation fee, charged one month in advance.

Once a subscription plan is activated, PortaBilling® immediately applies the subscription charges for the current billing period. At the end of the billing period, PortaBilling® applies the following month's charges in advance and adds them to the customer's invoice.

So if you assign a subscription plan to John on April 1st, PortaBilling® immediately applies both the \$10 activation fee and the \$30 periodic fee to John's balance.

At the end of April, PortaBilling® calculates all of John's usage charges for the current billing period (April) and applies the following month's subscription charges in advance (May). Thus, on May 1st John receives his invoice and sees charges for calls made in April plus the subscription charge for \$70 that covers:

- the \$10 activation fee,
- the \$30 periodic fee from April 1st to April 30th and
- the \$30 periodic fee from May 1st to May 31st.

On June 1st John receives his invoice and sees charges for calls made in May plus the \$30 subscription charge that covers from June 1st to June 30th.

An 'in advance' subscription plan can be activated at any time (e.g. on April 10th). In this case, the subscription fee is pro-rated so that the

customer is charged for the actual service usage (i.e. 20 days covering the period from April 10th to April 30th).

Let's say that John signs up for the service on April 10th. As described above, as soon as the subscription is activated, PortaBilling® applies the \$10 activation fee and the \$20 periodic fee (i.e. a daily fee [$\$30 \text{ monthly fee} / 30 \text{ days}$] * 20 days' usage) to John's balance.

At the end of April, PortaBilling® applies an 'in advance' subscription charge for May. Thus, on May 1st John receives his invoice with charges for calls made in April plus the \$60 subscription charge that covers:

- the \$10 activation fee,
- the \$20 periodic fee from April 10th to April 30th and
- the \$30 periodic fee from May 1st to May 31st.

You can charge your customers several billing periods in advance. In this case, PortaBilling® applies the full 'in advance' charge in the first billing period and then updates the subscription charges later on.

Have a closer look at this case. Let's say that you charge your customers 3 months in advance for their IP phone rental. Once Mark Roe signs up for the service on April 20th his subscription plan is activated. The \$20 charge is applied to Mark's balance (\$10 activation fee plus a \$10 periodic fee) for the period from April 20th to April 30th.

At the end of April, PortaBilling® calculates Mark's charges and applies the subscription charge for 3 months in advance. Therefore, on May 1st Mark receives his invoice and sees charges for calls made in April plus the subscription charge of \$110 that covers:

- the \$10 activation fee,
- the \$10 periodic fee from April 20th to April 30th,
- the \$30 periodic fee from May 1st to May 31st,
- the \$30 periodic fee from June 1st to June 30th and
- the \$30 periodic fee from July 1st to July 31st.

At the end of May, PortaBilling® detects that Mark's subscription charges cover only 2 months out of 3, since May is already paid for. Therefore, PortaBilling® applies a one-month subscription charge (\$30) which brings the total subscription charge to 3 months in advance once more.

On June 1st Mark receives his invoice and sees charges for calls made in May plus a subscription charge for \$30 that covers his IP phone rental for the period from August 1st to August 31st.

Subscriptions charged in advance to prepaid customers

In most cases, subscriptions charged in advance are assigned to prepaid customers. Therefore, an administrator must be sure that customers have sufficient funds to cover their subscriptions. When they don't, the customers' service is suspended.

The subscription plan configuration includes an option for automatically suspending prepaid customers' service when their funds are not sufficient to cover pending subscription charges.

This is how it works:

- Customers who have sufficient available funds have their subscription charges applied normally.
- Customers who have insufficient available funds are automatically suspended and no longer receive the service. Therefore, no subscription charges are generated. As soon as funds become available, the service is resumed and new charges are generated.

An administrator has the authority to suspend service due to insufficient funds for either a specific customer or for an entire customer class.

Let us consider the following example:

The customer has \$20 available funds and is assigned a \$30 monthly subscription for DID number usage charged a month in advance.

On November 1st the system calculates the subscription charges for the period between November 1st and November 30th (\$30) and compares that to the amount of available funds (\$20). The available funds are insufficient, thus the customer's service is immediately suspended.

On November 5th the customer makes a \$50 payment, so the available funds now total \$70.

The system recalculates the subscription charges. The customer did not use the service until November 5th, so the subscription charges are generated for the period starting November 5th (the day the customer became unsuspended and began using the services again) ending on November 30th (\$26). These charges are then compared to the amount of available funds (\$70), and applied. The amount of available funds then becomes \$44. Therefore the customer's service is reinstated. (Please see the *Invoicing* section in the [Unified PortaSwitch Handbook Collection](#) for a detailed description of the invoicing process.)

NOTE: After a suspension status is lifted from a customer, the system applies charges for the whole billing period (\$30) and generates a refund for any period during which service is suspended ($4 * \$30/30 = \4 where 4 is the number of days of suspension, \$30 is the total for subscription charges and 30 is the number of days in the billing

period). This refund is granted to the customer and two xDRs are created: one for the whole subscription charge (\$30) and another for the refund amount (-\$4).

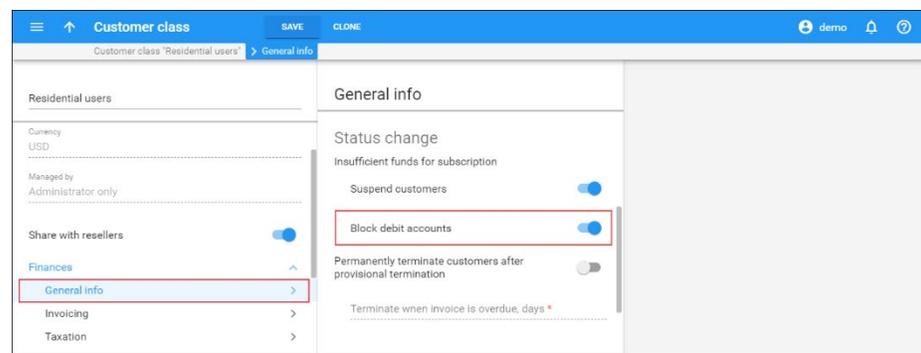
If the customer doesn't top up their balance in November, the following happens: PortaBilling® recalculates the customer's subscription charges each day to cover the remaining billing period and compares that with the funds available (\$20). Since the customer is not charged for the day they are suspended, the sum of subscription charges decreases by the daily fee (\$1). As soon as it equals the customer's available funds (\$20), PortaBilling® charges the customer and resumes their service.

PortaBilling® applies the charge for the whole billing period (\$30) and immediately issues a refund for 10 days of suspension. Thus, there are two xDRs produced for the customer: one xDR is for the whole subscription charge (\$30) and another is for the refund amount (-\$10).

Thus, prepaid customers can only consume services when they have sufficient available funds to cover their *next* subscription charges.

Subscriptions charged in advance to debit accounts

An administrator can automatically block services for debit accounts when funds are insufficient to cover pending subscription charges. As soon as an account once again has sufficient funds to pay for the subscription, the service is resumed.



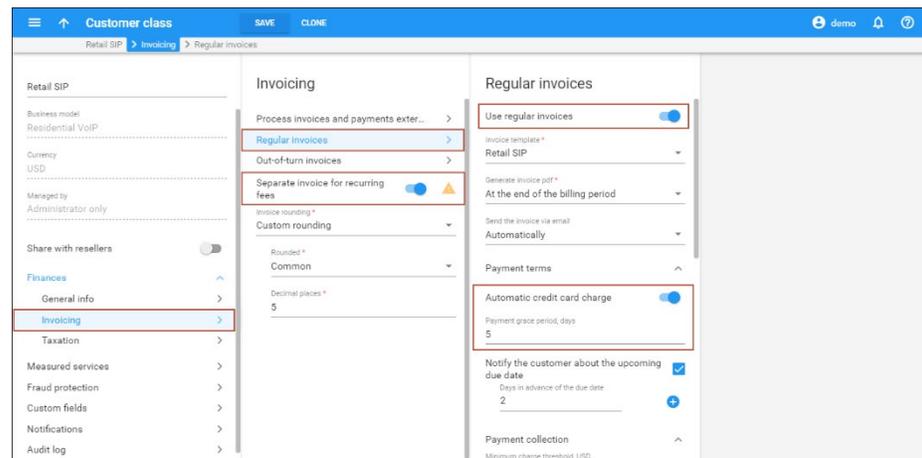
When the account is not charged for the days they do not use the service, PortaBilling® will issue a refund for the period during which the account was blocked.

Auto-charging for subscriptions

Another popular business model for applying subscription charges is to automatically deduct them from customers' credit cards. The service provider maintains a "service balance" for customers, that is, they may only spend their funds on service usage such as international calls, SMS, etc. This prevents situations in which the available funds are consumed by subscription charges.

To automatically charge customers for their subscriptions, configure a customer class as follows:

- Enable the **Separate Invoice for Recurring Fees** option;
- Enable the **Automatic credit card charge** option;
- Specify the invoice template;
- Specify the grace period in the **Payment is expected within _ days after invoice generation** option. For example, 5 days;
- Specify when to charge the credit card for the unpaid invoice in the **Automatic credit card charge after the due date, days** option. For example, 2 days after the invoice due date.



This is how it works:

The sum of subscription charges and calculated taxes for those charges is generated as a separate invoice at the end of the billing period. The invoice amount is automatically deducted from the customer's credit card, thus leaving the amount of available funds intact.

If the payment transaction is unsuccessful (e.g. if the credit card is expired or there are insufficient funds on it), the invoice remains unpaid and a corresponding notification is sent to the customer.

The customer has some grace period days for updating their credit card information. During the grace period the amount of available funds remains unchanged because a special correction xDR for the sum of the invoice is generated and inserted into the database. Thus, the customer can continue to use other services, such as make international calls.

NOTE: An invoice that contains subscription charges and taxes is paid automatically by a credit card only and cannot be paid manually via the web interface.

If credit card information is not updated within the grace period, the customer's service is suspended until the credit card information has been updated and the payment has gone through.

Consider the following example:

Prepaid customer John Doe has \$15 in available funds and signs up for a \$10 monthly subscription on April 1st. He is charged 10% in taxes for service usage.

During April, John Doe makes international calls that total \$6. His amount of available funds is now \$9.

On May 1st two invoices are produced for him:

- An invoice that covers the \$10 subscription charge and \$1 tax ($\$10 * 10\% = \1). This is an out-of-turn invoice in PortaBilling®.
- An invoice that covers \$6 for voice calls. Since this is a prepaid customer the invoice status is “do not pay.”

These invoices are mostly used for administrative accounting and are not usually sent to customers.

The sum of the out-of-turn invoice ($\$10 + \$1 = \$11$) is automatically deducted from John’s credit card so the status of the invoice is “paid.”

In May, John makes international calls totaling \$7, therefore his amount of available funds is now \$2. The invoices produced for him in June are as follows:

- An out-of-turn invoice that covers his \$10 subscription charge and a \$1 tax; and
- An invoice that covers \$7 for voice calls and has a “do not pay” status.

On May 31st John’s credit card expires, and consequently the out-of-turn invoice is not paid. The correction xDR for the sum of this invoice (\$-11) is generated and inserted into the database, leaving the amount of available funds unchanged (\$2).

On June 2nd John updates his credit card and the payment transaction for the out-of-turn invoice is successful. The invoice status becomes “paid” and the correction xDR is removed from the database.

If John did not update his credit card, his service would be suspended on June 3rd and terminated on June 18th.

With auto-charging for subscriptions, there is improved cash flow since there is pre-defined recurring revenue.

Customers are better able to control their available funds since these funds are only used for services.

Subscription refunds

A customer may close a subscription prior to the end date of the period that has already been charged for (e.g. the customer already paid for all of May but cancels his subscription on May 20th). Then a refund is issued for the ‘unused’ subscription time (e.g. 11 days).

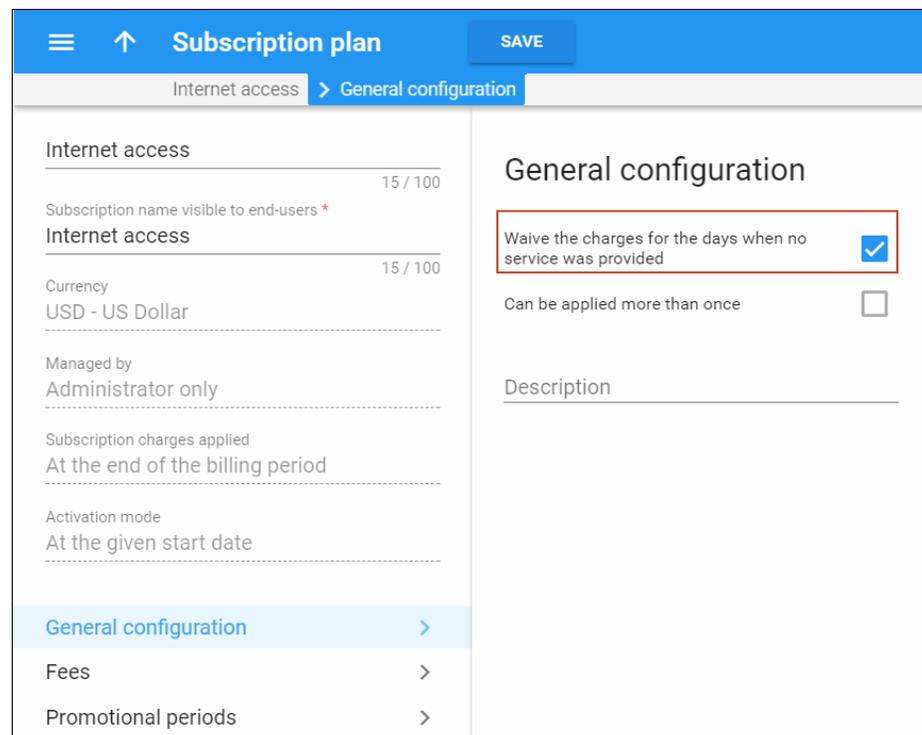
After the end date of the subscription plan the refund is calculated and immediately applied to the customer’s balance. The refund will then be reflected on the customer’s invoice at the end of the billing period.

Waiving the subscription fee when no service was provided

Service providers can choose whether or not to apply subscription charges for periods when their customers do not have access to services.

For example, customer John Doe is provided Internet service on a subscription-based model. The service provider decides not to charge John Doe for those periods when he does not have access to the service.

An administrator enables the **Waive the charges for the days when no service was provided** option for John Doe’s subscription.



The screenshot shows the 'Subscription plan' configuration page for 'Internet access'. The page is divided into two main sections: 'Internet access' and 'General configuration'. The 'Internet access' section includes fields for 'Subscription name visible to end-users' (set to 'Internet access'), 'Currency' (set to 'USD - US Dollar'), 'Managed by' (set to 'Administrator only'), 'Subscription charges applied' (set to 'At the end of the billing period'), and 'Activation mode' (set to 'At the given start date'). The 'General configuration' section includes a checkbox for 'Waive the charges for the days when no service was provided' which is checked, and a checkbox for 'Can be applied more than once' which is unchecked. There is also a 'Description' field which is currently empty. The page has a blue header with a 'SAVE' button and a breadcrumb trail showing 'Internet access > General configuration'.

On September 15th, John Doe reaches his credit limit. He cannot use the services until he tops up his balance. Due to financial problems, John Doe doesn't top up his balance until October 1st, so subscription charges for the period of September 15th – September 30th are not applied.

With the **Waive the charges for the days when no service was provided** option selected, subscription charges are withheld for periods when:

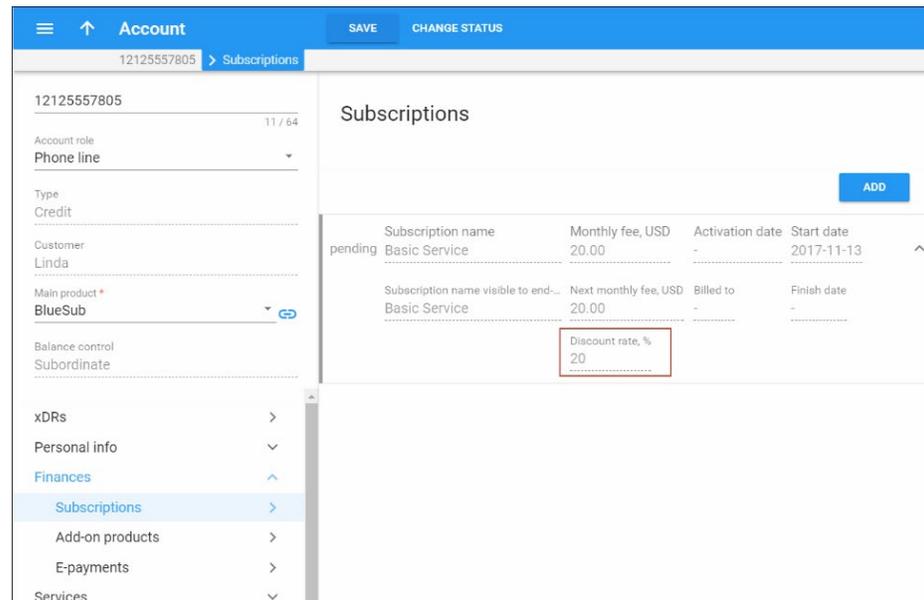
- A customer is suspended because of an overdue invoice.
- A customer or an individual account has reached the credit limit.
- A customer or an individual account has no available funds.
- A customer or an individual account is administratively blocked.
- An account is expired.

This approach increases overall customer satisfaction and encourages customers to continue using your services.

Subscription discounts

Discounts on subscriptions

Every subscription can have a custom discount rate (applied *only* to subscription's periodic fees), with the actual charged amount adjusted accordingly. Therefore, if a subscription plan defines the periodic fee as \$10.00, and this particular subscription is assigned a 20% discount rate, a charge of only \$8.00 is produced.



Use a 0 (zero) discount rate to specify “no discount” and ensure that charges are made according to the subscription plan definition. If a

subscription has an empty discount rate, the customer default discount rate is used (see below).

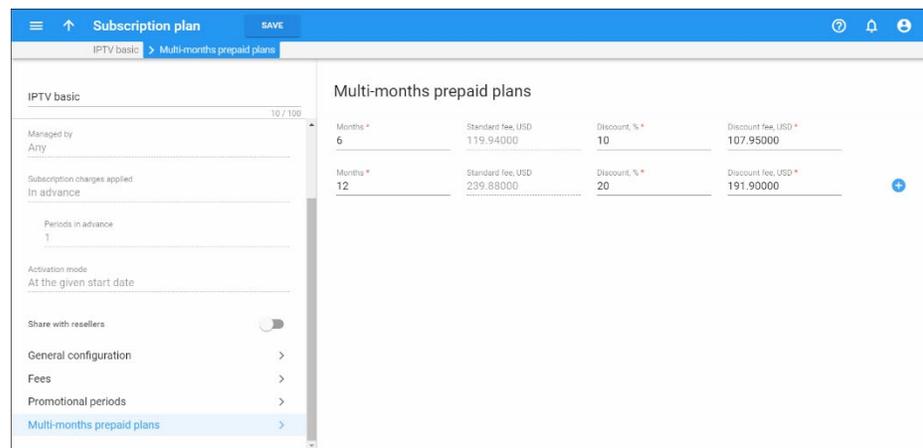
Customer discounts

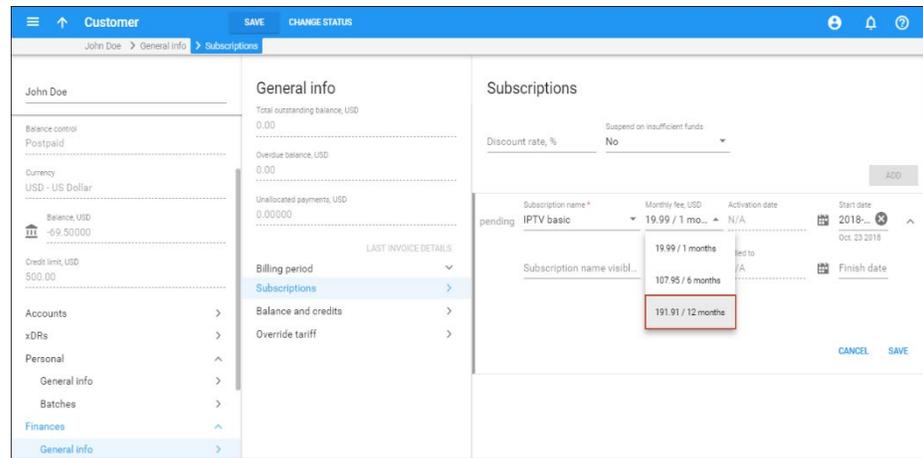
To simplify the management of different discount tiers (e.g. all your “silver” customers get 10% off the monthly fee for any of their subscriptions) you can directly assign a discount rate to a customer. So instead of creating multiple subscription plans with different periodic fee values, you will have a single set of subscription plans. You can then simply assign discount rates (10%, 20%, etc.) to customers. Any subscription (either directly applied to the customer or associated with one of the customer’s accounts) that does not have an explicitly assigned discount rate will then be charged using this discount.

Prepaid plans for subscriptions charged in advance

To improve your cash flow, you can encourage your customers to pay monies in advance by offering them different prepaid plans within subscriptions that are charged in advance (e.g. “Pre-pay for 6 months and get 20% off for Internet” or “Pay for 12 months and save \$100 for IPTV!”).

An administrator can set the number of months and choose a discount value, which will either be a percentage (e.g. 20%), or a fixed sum (e.g. \$20).





Then when assigning a subscription to a customer or an account, a corresponding prepaid plan can be selected.

Consider the following example:

In October 2015 customer John Doe signs up for the IPTV service with a \$19.99 / monthly fee charged in advance. He has two options for paying for the service:

- 10% discount for 6 months
 $(\$19.99 \times 6 - 10\% = \$119.94 - \$11.99 = \$107.95)$;
- 20% discount for 12 months
 $(\$19.99 \times 12 - 20\% = \$239.88 - \$47.98 = \$191.90)$.

John Doe chooses the 12-month payment plan so he is charged \$191.90 and saves \$47.98 (i.e. an xDR with this sum is created and inserted into the database). These charges are then reflected on his October invoice.

Moreover, he is not bothered with monthly payments until the end of the prepaid period (his invoices for the next 11 months will not include subscription charges since they have already been paid). The next subscription charge will occur in 12 months' time (i.e. in October 2016.)

Thus, customers have the option to choose how to pay for their services. When paying months in advance, they receive a discount on services and, as a result, save money.

ITSPs in turn improve their cash-flow, charging their customers for several billing periods at once and receiving larger sums as advance payment.

NOTE: In case of early closure of a subscription charged in advance, a refund will be issued for the "unused" subscription time not later than one day after the subscription closes.

Bundle promotions

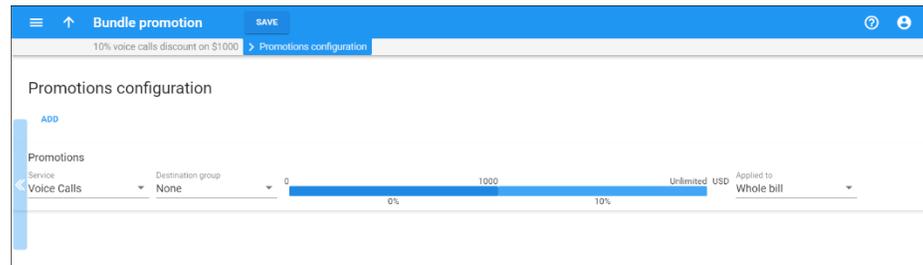
It is vital for any company to have a powerful marketing plan for attracting new clients and encouraging existing ones to spend more money on their services. The bundle promotions feature provides the following possibilities:

- to offer promotional discounts based on monthly amounts spent or volumes of the promoted service consumed;
- to define a mandatory minimum to be spent on a specific service and apply charges to customers for falling behind on that minimum.

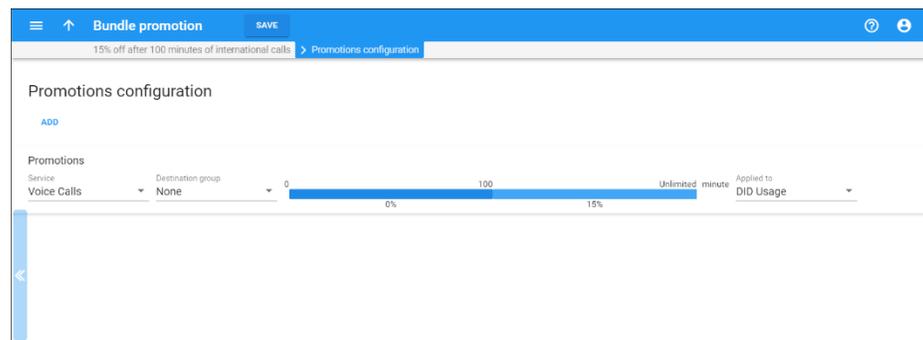
Promotional discounts

Key features of promotional discounts include the following:

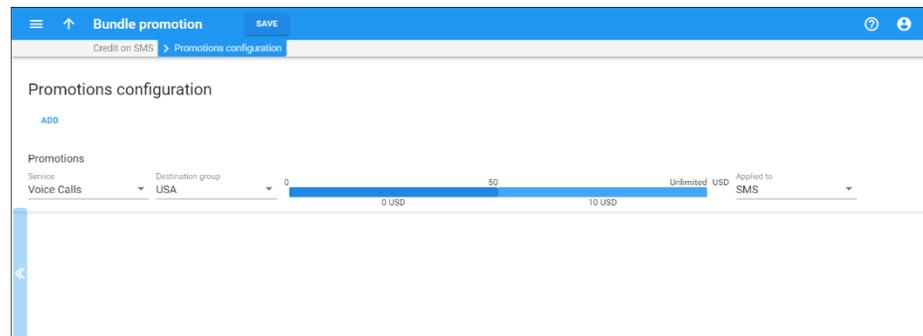
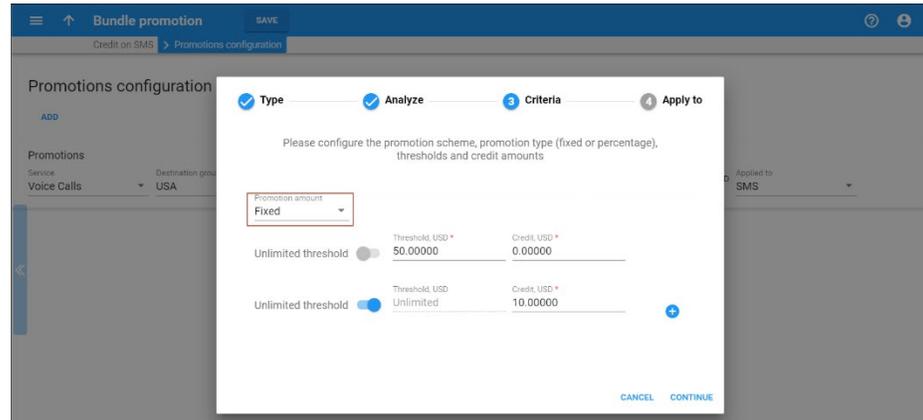
- Promotional discounts based on an amount of money spent for a service within the billing period. For example, if you wish to offer 10% off (of the invoice total) *if* the customer spends over \$1000 on voice calls (bringing the invoice total to \$900), you would create a corresponding destination group and define a single discount rule for it.



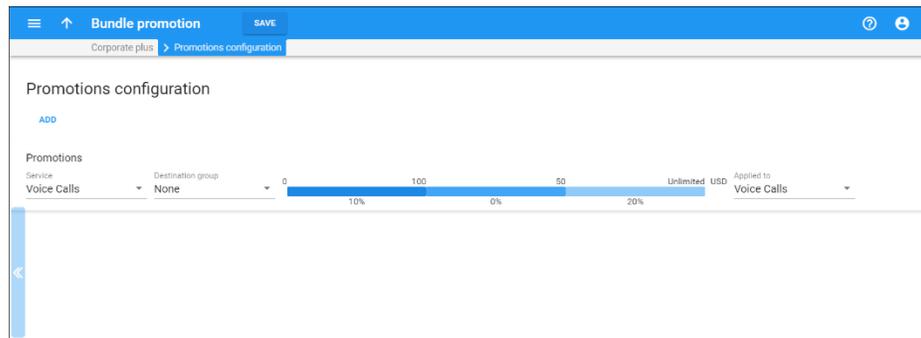
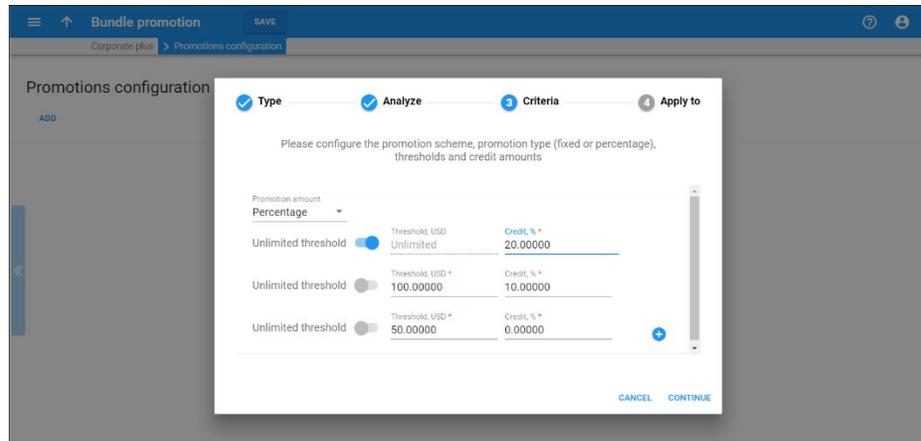
- Promotional discounts based on a volume of consumed service. For example, if you wish to provide a 15% discount on DID usage charges after making 100 minutes' worth of international calls, your discount rule will look like this: no discount is applied up to 100 minutes; a 15% discount is applied over 100 minutes.



- Promotional discounts given for actual service use of promoted services. For example, if you wish to provide \$10 off for SMSs *if* a customer spends more than \$50 on voice calls, define the promotional discount as a *fixed value*. Thus, a customer who has spent \$50 on calls and \$8 on SMSs, will receive an SMS discount of \$8 (of a possible discount of \$10).

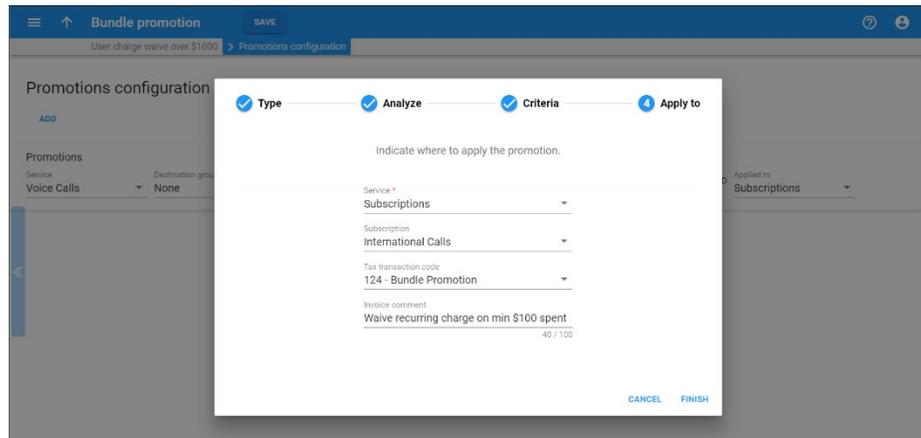


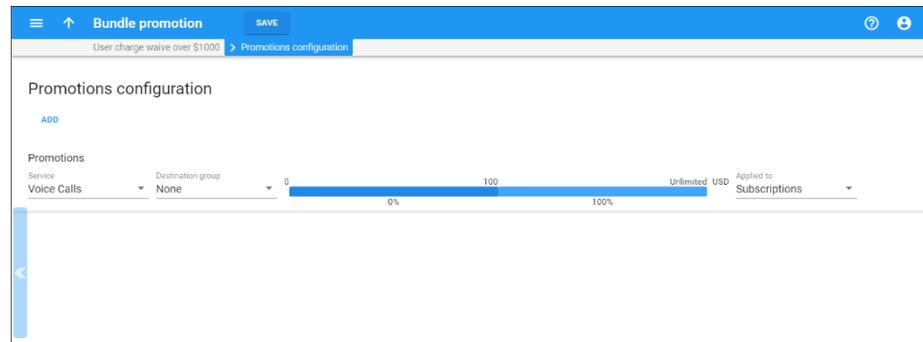
- The discount can either be a percentage of the amount charged in the 0 – 100 range (a 0% discount means “no discount applied,” while a 100% discount means “all charges waived”), or a fixed value.
- Multiple discount thresholds. For example, you can define a promotional discount rule for calls to Asia (thus only the amount of money spent on calls to Asian destinations will determine whether or not the discount should be applied). The discount will be applied to voice call service (cost of all calls within the billing period) so that: up to \$50 – no discount is applied, a 10% discount is applied for \$50 worth of calls and a 20% discount is applied for spending over \$100.



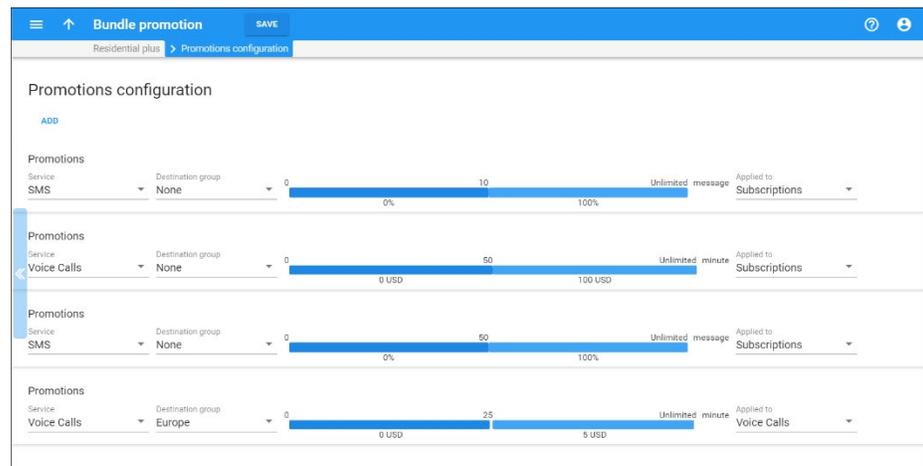
- Decrease a subscription fee or all subscription fees applied to the customer and his credit accounts by a defined percentage. For example, you can waive a customer’s monthly user charge if he spends at least \$100 on voice calls.

Note that if you select **Subscriptions** in the **Service** list and leave the **Subscription** list empty, the promotion will be applied to customer’s own subscriptions and *all* of his credit accounts’ subscriptions.





- Multiple promotions in the same promotion bundle. For example, your promotion bundle could state: “Spend \$10 dollars on SMSs and your subscription fee will be waived!” or “Spend \$50 either on calls or on SMSs and all your monthly fees will be waived,” or “Spend \$25 on calls to Europe and receive a 5% discount on all voice calls.”



- Promotional discounts based on charges for measured services. Let’s say you suggest a 15% promotional discount for international calls, for example, if the customer’s charges for extensions for that month exceed \$100. Then, you create a corresponding destination group and define a single discount rule for it.

Note that if you leave the **Destination Group** field empty, it means that all transactions related to this service will be summed up and analyzed for applying either a credit or a charge.

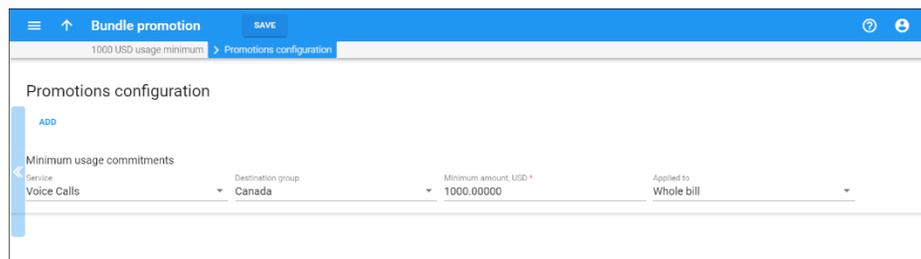
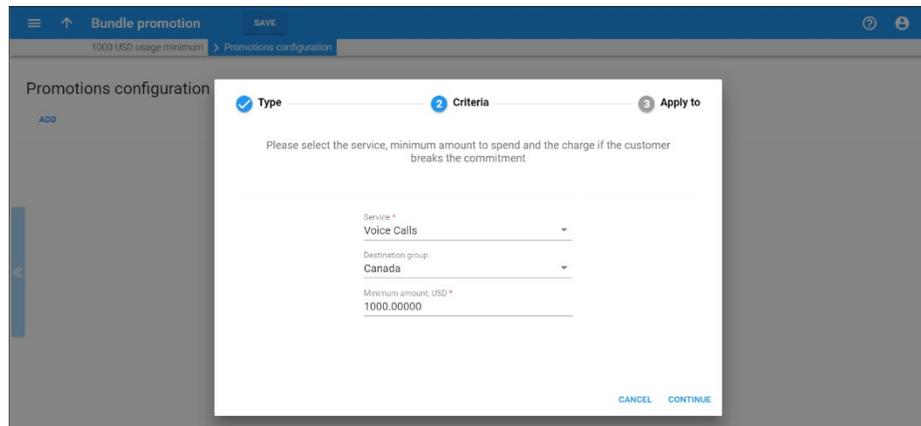
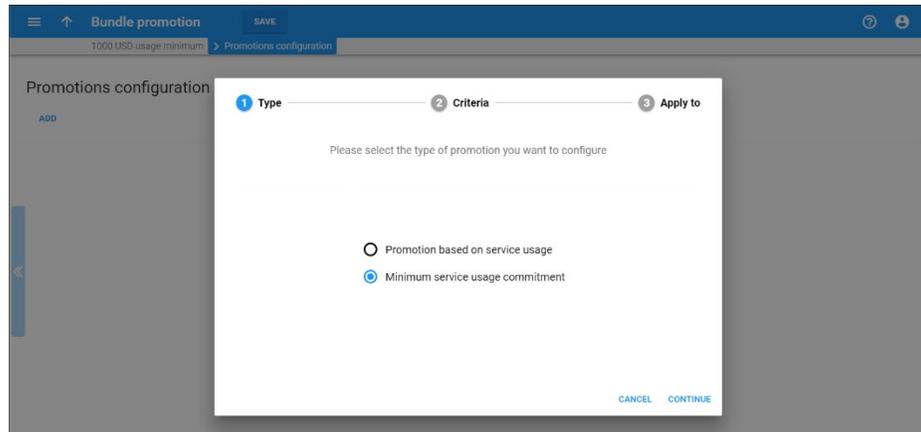
Minimum service usage commitment

You can define a mandatory minimum for customers who commit to using a specific quantity of service. The mandatory minimum ensures that customers’ service charges are always equal to or higher than a predefined amount.

You can apply a minimum service usage commitment either to the same service or to the invoice total.

Consider the following example:

A customer signs a contract guaranteeing great rates but he is obligated to spend \$1000 minimum per month on voice calls. In fact, the customer only spends \$800, so as a result, the deficient \$200 amount is added to his invoice so that it equals \$1000.



This approach means that your revenue stream is more predictable. Therefore, you can plan your expenses more efficiently and become more financially stable.

Bundle promotion calculation

A bundle promotion is applied at the end of a customer's billing period – after all customer service charges and balance adjustments have been processed but before taxes are calculated – and is reflected in (a) separate line(s) on the invoice.

If a bundle promotion is applied to a particular service (e.g. Conferencing), it is calculated only for the sum of charges pertaining to this service usage. If a bundle promotion is applied to the whole bill, PortaBilling® calculates it for the sum of charges for voice calls, subscription charges and balance adjustment operations (manual credits) that were performed within this billing period.

Let's look at the following example:

You apply the promotion: “10% off if the customer spends over \$1000 on voice calls,” (see the first example above) and define that the promotion applies to the total invoice amount.

At the end of the month, the customer has spent \$1200 on phone calls.

When the customer's billing period closes, PortaBilling® calculates the bundle promotion as $\$1,200 * 10\% = \120 . An extra xDR is created and the customer sees a line with a \$120 **credit** on his invoice. The description for this transaction is **10% (\$1,200)**, making it is easy to understand what it means. The invoice total is then \$1,080.

The next example illustrates how promotional discounts are applied for volumes of consumed services:

Let's say your promotion states: “Get \$20 off your invoice total for making over 100 minutes' worth of calls per month – and \$30 off for making over 200 minutes' worth of calls per month!” Within a month, a user has made 130 minutes' worth of calls with a \$15 cost in total. Thus, his invoice total is adjusted by \$15.

The differences between bundle promotions and volume discounts

Unlike volume-based discounts that apply to the same service, bundle promotions allows you to encourage customers to use one service to obtain a discount for another service, or increase or even decrease a customer's invoice total, based on how much money has been spent on these services during the billing period.

	Bundle Promotions	Volume Discounts
Scope	Discount can be applied to a different service.	Discount applied to the same service.
Discount base	Discount base for credit amount is calculated based on the total service usage within this period.	Discount is applied only to service consumption after the threshold is reached.
Discount amount	Can be a fixed amount or calculated as % of the sum of charged amounts.	Calculated as a % reduction of the rate.
Applied to	Discount can be applied to any service or subscription charges.	Discount can only be applied to a service.
Discount applied	At the end of the billing period.	In real-time (when the service is consumed).
Presentation	Separate xDR on the invoice.	Incorporated in the original xDR.

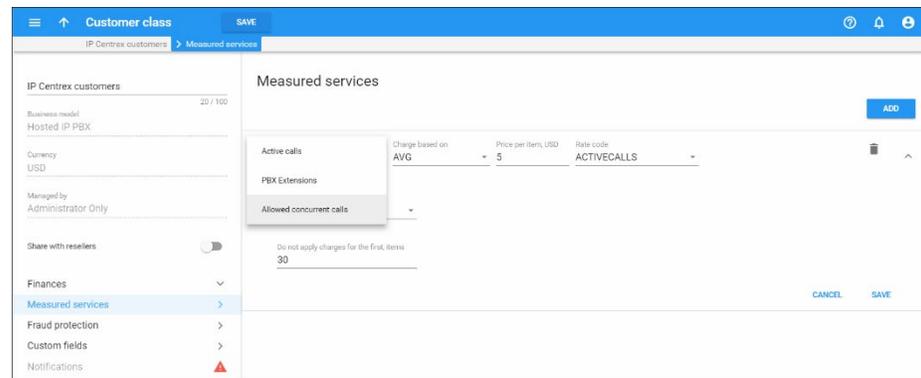
Billing based on measured resource consumption

PortaBilling®'s ability to measure allocated resources allows service providers to implement this business model. The aggregated value of a measured resource can be used for statistics and for applying charges to customers, as well.

The parameters that can be measured and billed for are as follows:

- The number of concurrent calls allowed (SIP trunks). This is the number of simultaneous outgoing calls allowed for a particular customer. The number of simultaneous outgoing calls is defined by an administrator as part of the customer service configuration for all customer accounts or for this customer's individual site (in the **Limit simultaneous calls** service feature);
- The number of active calls. This is the *actual* number of connected calls at a specific moment in time;
- The number of IP Centrex extensions. This is the number of extensions that a customer defines within their IP Centrex environment.

This list is provided by default; however, it is easily extensible and in future releases will include storage space for call recordings, number of IP Centrex phone lines and other parameters.



Resource consumption is measured periodically at different time intervals for different parameters. Thus, the measurement frequency is as follows:

- The number of active calls is measured every 5 minutes.
- The number of permitted concurrent calls is measured daily.
- The number of IP Centrex extensions is measured daily.

This measured data is aggregated and used for displaying consumed resource statistics and for calculating charges.

Charging for measured resource consumption

Charges for resource usage are calculated based on the following criteria:

- the minimum,
- maximum, or
- average resource consumption.

These charges are applied to customers or resellers as part of closing their billing period (an xDR is created and added to the invoice on a separate line). Criteria for charges are defined for a customer class and if necessary, can be overridden for individual customers or resellers.

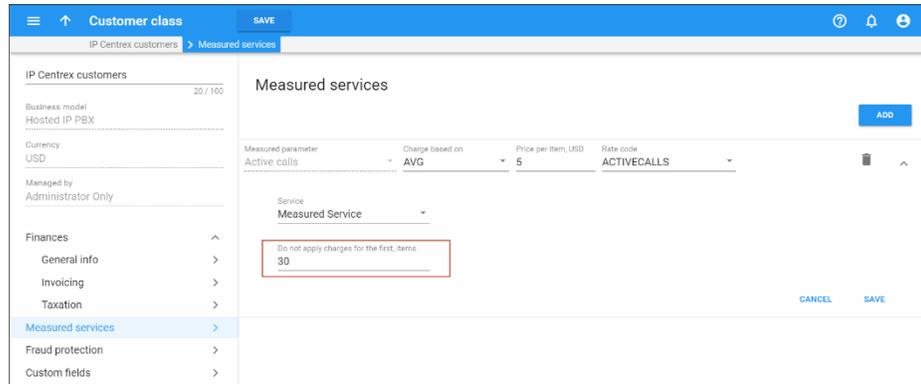
Consider the following example:

The company EasyCall Ltd. has a limit of 30 simultaneous outgoing calls according to their contract with their service provider. Yet this month they are organizing a calling campaign to their customers to inform them about their legal address change, so they have requested that the service provider increase their simultaneous calls limit to 100.

After a week, their calling campaign is finished and their threshold of simultaneous calls is limited once again to 30. At the end of the month, the simultaneous calls for the current billing period are averaged up and calculated as 46.3, which is rounded up to 47.

Thus, EasyCall Ltd. is charged for 47 simultaneous calls.

As a rule, a customer’s service bundle already includes some amount of allocated resources covered by a subscription fee. This amount is specified in the **Do not apply charges for the first items** field.



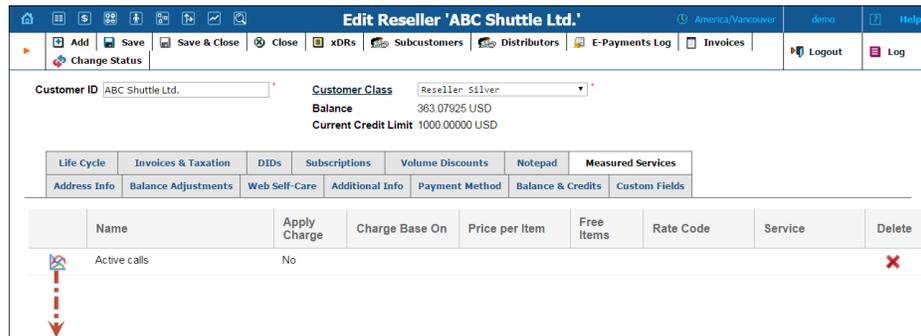
If customers consume more of an allocated resource than is defined in their service bundle, they are charged extra for additional resource usage. Thus, if a customer’s service bundle includes 30 active calls in their \$30 monthly subscription and the average number of active calls for this customer within the current billing period is 50, they are only charged for 20 active calls (50 – 30 = 20).

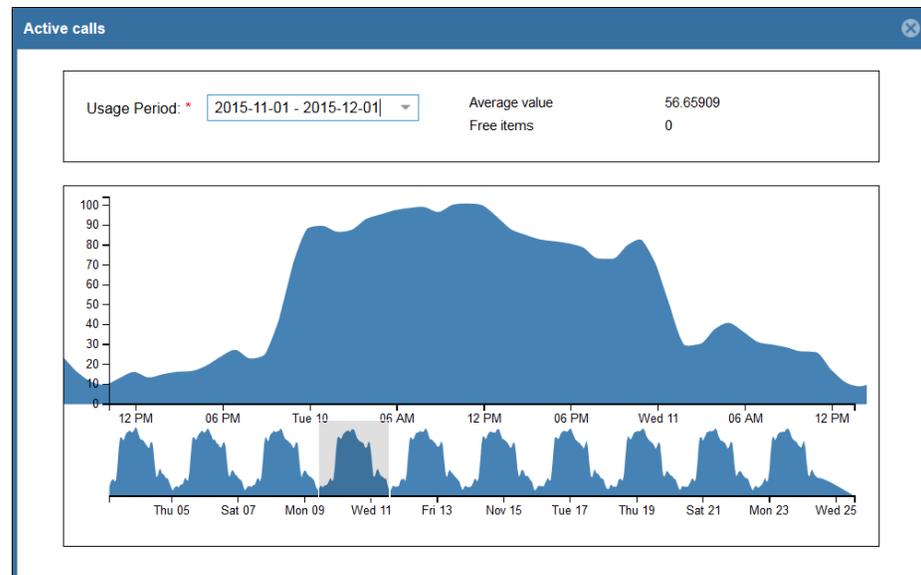
Measured resources statistics

The measured resources statistics are available as a graph on administrator, reseller and customer self-care interfaces. Thus, administrators and resellers can browse their direct customers’ and resellers’ statistics on the **Measured Services** tab of the **Edit Customer** page. Resellers can also browse their own statistics on the **Measured Services** tab of the **My Company Info** page.

Statistics for customers are provided on the **Billing Information** tab of the customer self-care interface

To browse a customer’s statistics, the administrator clicks the  icon next to the necessary resource metric.





The statistics window consists of two graphs. The bottom graph displays resource usage for the previous three months; the upper graph displays a detailed view of resource usage for a particular time interval selected from the three-month period, below.

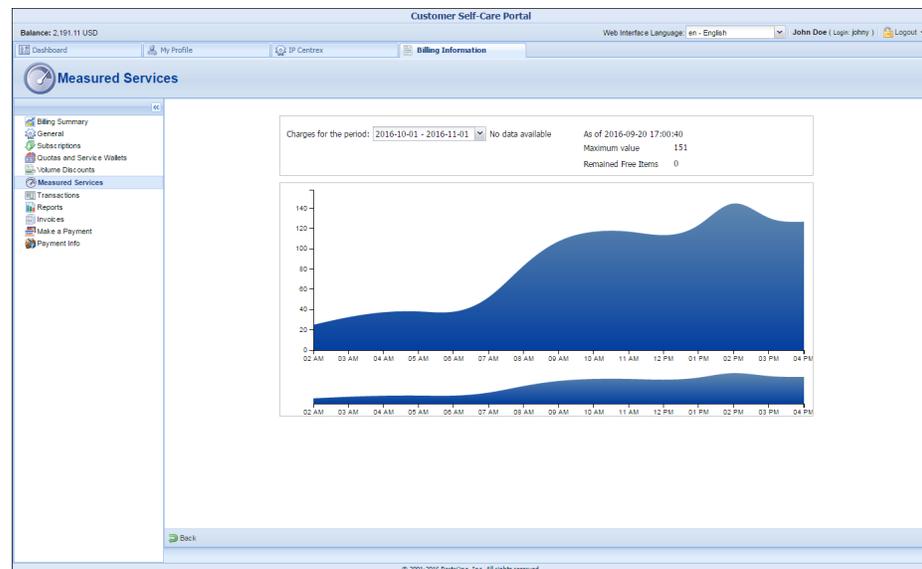
The topmost part of the window provides information about charges for selected resource usage as follows:

- **Charges for the period** – the customer's billing period. Any billing period within the xDR storage time can be selected. For details about storing xDRs please refer to the **xDR Cleanup Procedure** chapter.
- **Amount** – if a customer is charged for resource consumption, those charges are reflected in the selected billing period.

NOTE: The charges for the current billing period will not be displayed since this billing period is not yet closed.

- **Minimum / Maximum / Average Value** – the measured resource value for the current billing period, depending on the charging criteria defined for this customer.
- **Remaining Free Items** – the number of available allocated resources that are included in the service bundle for the current billing period.

The same information is provided for customers on the customer self-care interface by clicking the  icon.



Thereby the measurement of allocated resources serves as an additional system control tool and at the same time, charging for actual resource consumption improves your revenue flow.

Charging customers with hierarchy for the number of extensions

Service providers have advanced flexibility for charging customers with distributed infrastructure for the number of their extensions. (For a detailed description on how to organize companies that have several functional departments in PortaBilling® please refer to the **Customer Hierarchy** section). In other words, administrators can charge the main office for all and / or some of the extensions used, or charge each office individually.

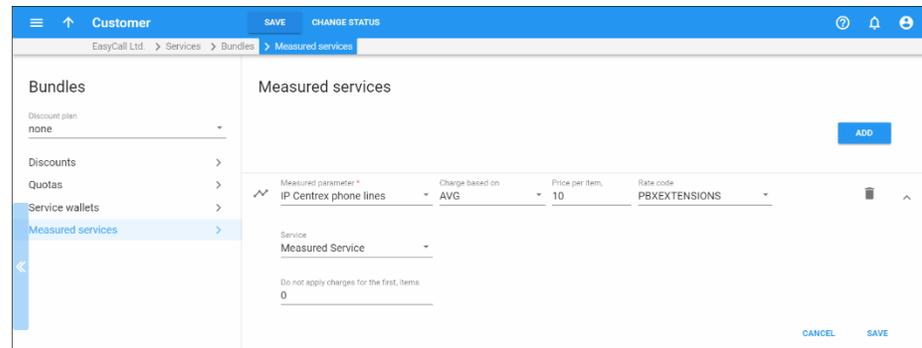
Let's say, for example, that customer EasyCall Ltd. has a main office and 2 branch offices (Branch 1 and Branch 2). Each office has several extensions:

- The main office has 2,
- Branch A has 3, and
- Branch B has 4.

Since there is a \$10 charge for each extension, an administrator can configure charging in one of the following ways:

1. The main office is charged for all extensions:

An administrator configures the measured services functionality for the main office. At the end of the billing period, the main office is charged \$90.



2. The main and branch offices are charged individually:

An administrator configures the measured services functionality for each office. At the end of the billing period the main office is charged \$20, Branch A – \$30, and Branch B – \$40. Those charges are added to each office's invoice.

Administrators can also configure PortaBilling® to charge the main office for both their and some of the branch office's extensions. Let's say that the customer wants the charges for Branch A's extensions to be added to the main office's invoice, and that Branch B be charged separately for their extensions. The administrator then configures the measured services functionality for just the two offices: the main one and Branch B's office.

In this case, the main office is charged \$50 (\$20 for its own extensions plus \$30 for Branch A's extensions) and Branch B is charged \$40. Those charges are added to the main office and Branch B's invoices at the end of the billing period.

3. Only branch offices are charged for extensions:

An administrator configures the measured services functionality for each of the branch offices. Then each branch office is charged individually, so Branch A is charged \$30 and Branch B is charged \$40. Those charges are added to each branch office's invoice at the end of the billing period. The main office is not charged.

This provides a more flexible billing structure for companies with a distributed infrastructure.

5. Voice services

Rate upload

The Rate Upload functionality is a powerful and flexible tool that allows you to save time while uploading new rates to the tariffs. Let's take a practical example of the rate upload procedure: vendors send you their rates and you have to enter this data into the system. Obviously, regardless of the particular numbering plans, formats and other options, this data has to be in the unified format used in PortaBilling®. Some vendors use Microsoft Excel, others text-based formats (for example, .csv). Even if the same format is used, there will be different ways of positioning and formatting the data. For example, vendor A might have the phone prefix for the Czech Republic in the first column in his data file, entered as **420**. Vendor B might have it in column number three, entered as **011420**.

Of course, it is possible to write an external tool which would convert data files from these different formats into a common one. However, this would require some extra steps in data processing, so that every time you upload rates from a vendor you will have to perform a conversion first. In addition, this will probably involve extra cost, since not all users are capable of writing the required data conversion application.

This is why PortaBilling® has chosen a different approach. Instead of making your data fit PortaBilling®, you can make PortaBilling® understand and process your data as is! The PortaBilling® rate upload wizard and templates allow you to do the following:

- Upload your .xls or .csv rate file to the server and immediately see the result of file processing (in order to verify that the file was recognized correctly, and to adjust parameters such as the field delimiter).
- Using a drag-and-drop interface, you can place markers on columns in the file to identify where particular groups of data are located in it (e.g. the destination prefix in the third column, and the price in the fifth column). In fact, the rate upload wizard will try to automatically recognize data elements, so most of the time you will only need to confirm the proposed column assignment.
- You can adjust the format which each data element is in (for example, whether the phone prefix is given as **011420** or as **420**). You can also apply post-processing rules (analogous to formulas in spreadsheet editors).
- The upload wizard will process all the rate information and allow you to compare new rates with the existing rates and make adjustments if necessary.
- If there are new destination prefixes in the file (not in the database yet), they can be created automatically during rate upload. Also,

you can assign new prefixes into specific destination groups during the upload process.

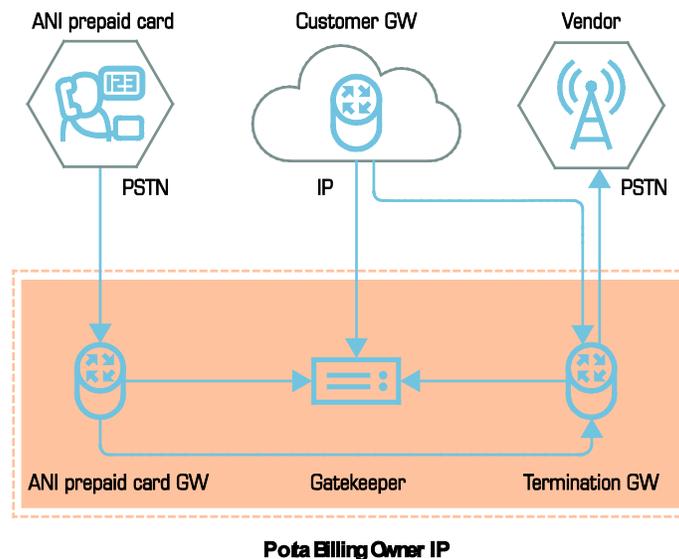
- Once the new rates are uploaded into the system, PortaBilling® saves all of your changes as a template associated with this tariff. Thus, the next time you receive a new file from the vendor and need to upload it, everything is ready and no special actions are required.

With Rate Upload functionality you save time while uploading new rates to the tariffs. This is achieved by saving all of the steps performed as a template associated with this tariff that can be used later on during subsequent rate uploads. Fewer errors occur while uploading rates with the help of the Rate Upload wizard since all processing and conversion is visible and controlled by the user.

xDR calculation

PortaBilling® can collect information about the call via various network elements (e.g. PortaSIP® and the media gateway), combine the individual pieces and produce a single xDR record.

Collecting data from the VoIP network



Gateways in the PortaBilling® owner network are “trusted” gateways. These gateways **MUST** be listed in the Nodes directory and be a RADIUS or Diameter client.

Gateways which belong to your customers or vendors should not be created as nodes. Instead, a remote GW which belongs to your customer will be represented in the system as an account.

If the termination GW transmits an accounting record (connection matched), the User-Name must be verified to see if a trusted GW (Node IP=User-Name) is calling. If so, then the billing engine must wait for an accounting record from the trusted GW and use the User-Name from a trusted GW accounting record. If a trusted GW is not found, the call will be charged immediately to an applicable account. If an account is not found, a “security breach” mail alert will be sent.

Call legs

Billing for telephony services is done based on the information provided by the elements of your VoIP network: gateways, SIP server, and so on. This information is supplied to the billing engine as accounting records. Usually the gateway treats incoming and outgoing calls as two separate entities, i.e. **call legs**. Each of the call legs produces a separate accounting record, so that, for instance, billing engine receives 4 accounting requests in case of a “classic” VoIP call arriving from PSTN to gateway A, traveling over the Internet to gateway B, and terminating in the PSTN network once again. Of course there must be some way of knowing that these four records are actually part of the same logical call. Therefore, each of them includes a special attribute that carries a unique ID for the call. This ID is used by the billing engine to add the call legs together, and is also extremely useful for troubleshooting. For H323 calls, this attribute is called `h323-conf-id`, while for SIP calls it is `call-id` (however, for backward compatibility Cisco gateways and PortaSwitch® assign a `h323-conf-id` for every SIP call).

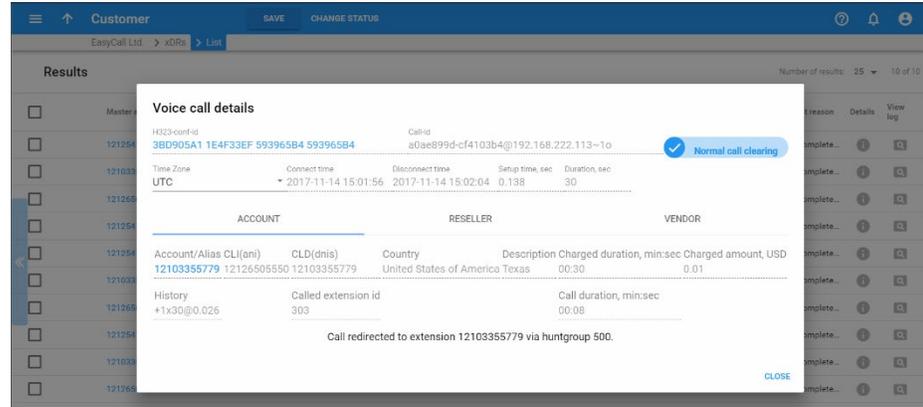
Extended xDRs for IP Centrex calls

Owners of IP Centrex environments may wish to obtain information about the actions of particular extensions during “complex” calls, i.e. calls that consist of several call legs such as call transfer, call forwarding, etc. For example, you might want to know which extension transferred the call or for how long an operator actually spoke, so that this information can be used in statistics reports.

PortaBilling® compliments xDRs for complex IP Centrex calls with auxiliary call details such as:

- Which huntgroup a call came from.
- What the last extension ID was that forwarded or transferred the call.
- The actual call duration, i.e. the length of time an operator spoke in the call.
- The caller identity from where the call originated.

All of these details are available on the administrator web interface as part of the xDR history for a customer or their accounts. To browse auxiliary call details, click the **Details** icon next to the required xDR.



The **Voice call details** dialog window displays:

- The extension number to which the call was redirected.
- **Via** – The extension or huntgroup number that dispatched the call.
- **Call duration** – The time during which the owner of the account being billed was occupied.
- **Caller extension ID** – The extension number where the call originated from.
- **Called extension ID** – The extension number that accepted the call.
- **Calling party identity** – The actual caller identity where the call originated from.

Auxiliary call details are created for the following IP Centrex calls:

- Calls to voicemail (direct and missed calls).
- Forwarded / follow-me calls.
- Call pick-up.
- Transferred calls (via blind and attended transfer).

This list is provided out of the box. However, it is easily scalable and in future releases will be extended to cover more call scenarios such as call queues, etc.

External custom reports

The auxiliary call details are defined by the following xDR fields:

- `via_hg` – Information about the huntgroup that dispatched the call.
- `via_ext` – Information about the last extension ID from where the call was forwarded or transferred.

NOTE: These flags appear when a call is processed within the IP Centrex environment.

- `media_duration` – Information about the actual time the owner of the account to be billed was speaking during the call.
- `accli` – Information about the actual caller identity within the IP Centrex environment from where the call originated. This identity may differ from the one delivered by the IP PBX (e.g. the IP PBX main account or the caller’s phone in case of forwarded calls) and displayed on the device of the called party.

These fields are added to the `XDR_Auxiliary_Types` table in the database while their values are stored in a separate `XDR_Accounts_Aux` table. The administrator can retrieve the information by accessing the database and can formulate a statistics report using a third-party application such as the Crystal Report tool. Please refer to the [Custom Reports](#) section for details of the custom report creation.

With extended xDRs you receive more sophisticated statistics data which helps you improve the IP Centrex management. The data may also serve as an additional tool for accounting.

Understanding Charged duration and Call duration in xDRs

When viewing **Voice call details** in customer and account xDRs, you see two fields: **Charged duration** and **Call duration**. What are they and what’s the difference between the two?

Voice call details

H323-conf-id
C03EEFB4 86428E68 1EE28707 1EE28707

Time Zone
Europe/Kiev (user)

Call-id
341425f-9b165310@192.168.246.191

Connect time
2019-12-25 04:51:09

Disconnect time
2019-12-25 04:51:20

Setup time, sec
0.437

Duration, sec
11

✔ Normal completed call
Normal call clearing

ACCOUNT		RESELLER		VENDOR		
Account/Alias	CLI(ani)	CLD(dnis)	Country	Description	Charged duration, min:sec	Charged amount
12014444444	12012233333	12014444444	Not set	Not set	00:11	0.000... USD
History		Call duration, min:sec				
+11x1@0.0		00:11				

[DOWNLOAD CALL RECORD](#)

[CLOSE](#)

Charged duration is the total session duration. It is recorded in the `charged_quantity` xDR field. PortaBilling® uses this value to calculate charges for the corresponding account. (The account to bill for the session is stored in the `User-Name` attribute).

Call duration is the time the user actually speaks. It is recorded in the `media_duration` xDR field and is mainly used for statistics purposes.

In direct calls between users, these two field values may be equal. In complex calls such as call transfer, call forwarding, call pickup and calls that involve an IVR, they may differ. Let's have a look at some examples:

Example 1. Call forwarding

An external caller Mary calls Bob. Bob has follow-me configured so the call is forwarded to Carol. Mary and Carol speak for 5 minutes. After the call ends, PortaBilling® produces the xDRs for the call participants: Bob is charged for the incoming call from Mary. Since Bob forwarded the call to Carol and was not physically speaking but is the one to pay for this call, his xDR details include:

- charged duration of 5 minutes, and
- call duration of 0 minutes.

Carol is charged for the incoming call from Bob. Her xDR details include:

- charged duration of 5 minutes, and
- call duration of 5 minutes.

Example 2. Blind transfer

Alice calls Bob. After 1 minute of conversation, Bob transfers the call to Carol. Carol answers and Bob is disconnected. Alice and Carol speak for 10 minutes.

After the call is over, PortaBilling® produces the xDRs for call participants:

- For Alice: the xDR for the outgoing call to Bob. The xDR has the charge for the whole call, which includes the call to Bob (for 1 minute) and the call to Carol (for 10 minutes). Thus, in xDR details for Alice we see charged duration and call duration equal 11 minutes.
- For Bob: the xDR for the incoming call from Alice. The charged duration and call duration are 1 minute.
- For Bob: the xDR for the redirected call to Carol. The charged duration is 10 minutes (Bob's account is billed for the transfer) but the call duration is 0, since Bob did not participate in this conversation.
- For Carol: for the incoming call from Alice. The charged duration and call duration are 10 minutes.

Example 3. Call to a call queue via the auto-attendant

Alice calls EasyCall. She dials 12065558899 – the auto-attendant for EasyCall. She presses 1 and is redirected to the Sales queue. After 2

minutes she's connected to Peter, the sales agent. They speak for 7 minutes.

After the call ends, PortaBilling® produces these xDRs:

- For Alice – for the outgoing call to Pete. The xDR includes the charged duration and the call duration of 9 minutes (2 minutes for waiting in the queue and 7 minutes for the conversation).
- For the auto-attendant account (12065558899) – for the redirected call from Alice to Pete. The charged duration is 9 minutes because the auto-attendant account is billed for the transfer to Pete. The call duration is 0 because the auto-attendant does not participate in the conversation.
- For Pete – for the incoming call from the auto-attendant. The charged duration and the call duration equal 7 minutes.

xDR recalculation

Mistakes can happen – someone sends you the wrong pricelist, or your administrator simply clicks the wrong button, resulting in incorrect charges in the database. How this can be corrected before an angry customer calls you?

PortaBilling® provides two methods for fixing operator errors:

- RADIUS re-feed
- xDR re-rating

RADIUS re-feed

Every accounting request received by the billing engine is stored in a separate file on the billing server. A special command-line utility supplied with PortaBilling® can be used to extract all RADIUS requests for a specific time interval (starting from the time when the error was made until the time it was fixed) and resend them to the billing server. Each request is tagged, so that a special processing mode is invoked.

- Any existing xDRs for this session are erased from the database, and balances are rolled back (as if the previous charge attempt had never happened). This also includes turning back any volume discount counters that were modified.
- The session is charged according to the new (correct) billing configuration.

Re-feed is a very powerful tool, and can fix any type of billing error (including missing connection to vendor, incorrect product rating table, or wrong rate). However, it must be used with caution, as it involves processing large volumes of data. Customers are advised to contact the PortaOne support team, who can provide further assistance in running the re-feed.

xDR re-rating

This is a lightweight and easy-to-use alternative to re-feed, designed to fix the most common problem: incorrect pricing information entered into a tariff. It is operated completely from the PortaBilling® web interface. The administrator can choose to run an xDR re-rating for either all customers or just a specific one, then specify the “original” (incorrect) tariff and correct tariff, and then narrow down the set of xDRs to be processed by time interval. Additionally, the administrator can specify that an override tariff be included in the xDR re-rating process.

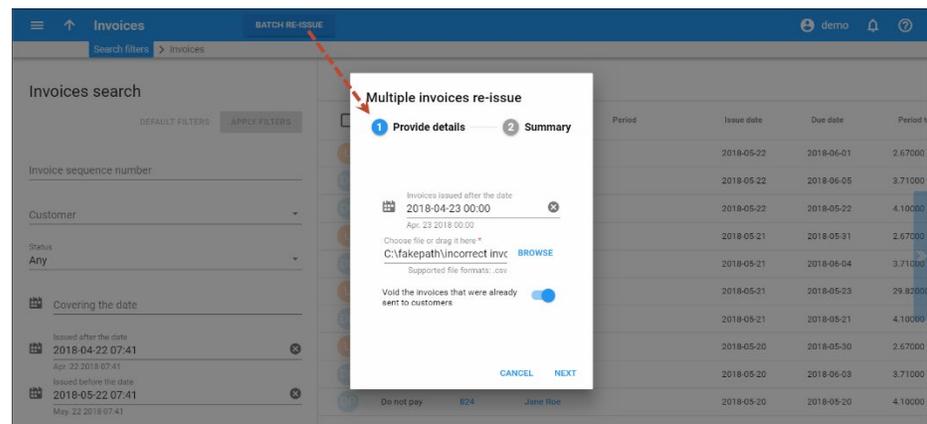
The PortaBilling® xDR re-rating feature can perform re-rating in two different modes. You can change the xDR re-rating mode per environment on the Configuration server (the **DeltaMode** option).

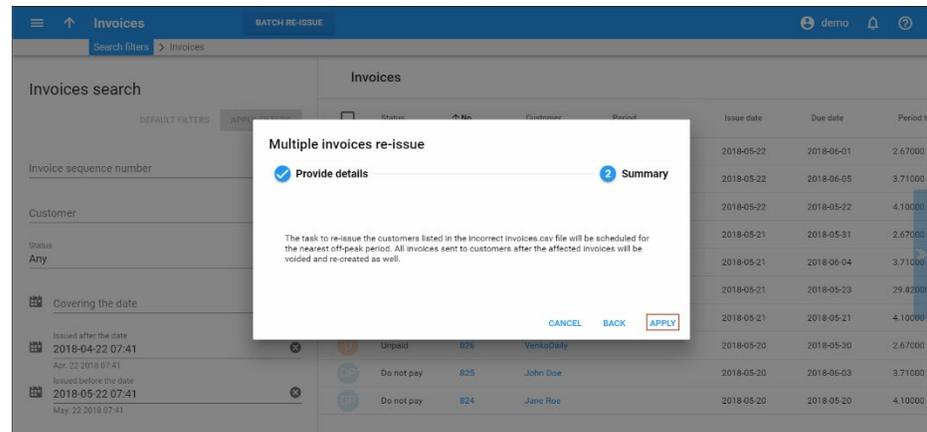
Correction mode

The re-rating task starts in the background, and during its execution:

- The script locates all xDRs produced according to the incorrect tariff (and matching time interval and customer criteria).
- For every xDR, the script calculates the charge according to the new tariff.
- The xDRs are then updated with the correct charged amount, and the balance is adjusted by the difference between the new and old charged amount.

However, the previous incorrect invoices are not automatically voided upon completion of the re-rating task. Therefore, it is necessary to void those previous incorrect invoices (that have already been issued and sent to customers) and produce new ones. This is done by starting the invoice re-issue process on the **Invoices** page.





NOTE: If the incorrect invoices have not been sent to customers, they will be recalculated without voiding.

Delta mode

If the xDRs produced according to the incorrect tariff fall within a previous (closed) billing period, then “corrective” xDRs are added to the current billing period and the original xDRs won’t change. These new xDRs contain the difference (delta) between the new (correct) and old (incorrect) charged amounts. Therefore, the previous invoices are not changed, and instead, a new updated invoice (that covers the current billing period) will include all of the corrected amounts.

In Delta mode, new xDRs will have the same service and timestamps as in the originals.

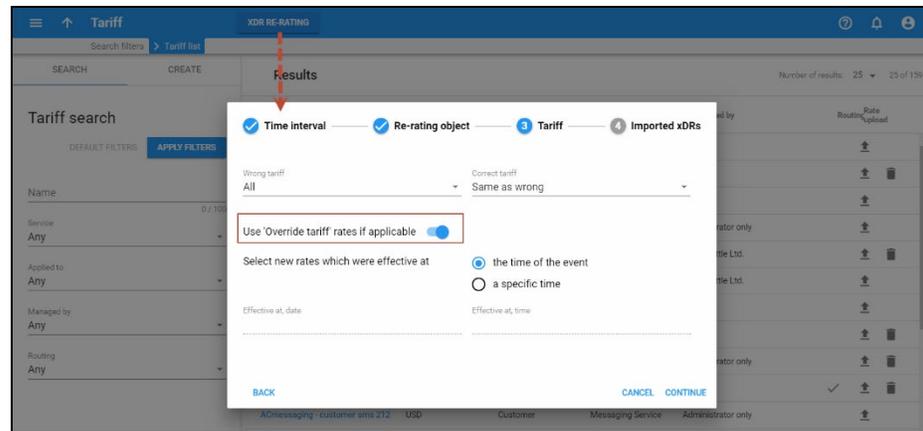
Re-rating and volume discount counters

If volume discounts are used when calculating call charges, it is no longer possible to treat such calls separately from others, since the way one call is charged affects all other calls made subsequently. For instance, if a call is charged \$5, this is the amount added to the volume discount counter. If the charged amount is then changed to \$1 during re-rating, this will affect all other calls in the same destination group, since a different discount would now apply to them.

To overcome possible confusion when volume discount counters are involved in re-rating, the re-rating procedure should be run within a specific period in the past (starting from the time when the error occurred, ending with time when it was fixed.) All discount counters will be rolled back before recalculation actually starts, and then updated with each re-rated call.

xDR re-rating and override tariffs

For a customer who has an already configured override tariff, instruct PortaBilling® to use it during the xDR re-rating. To do this, enable the **Use Override Tariff** option.



Then during the xDR re-rating, PortaBilling® will check the customer override tariff and:

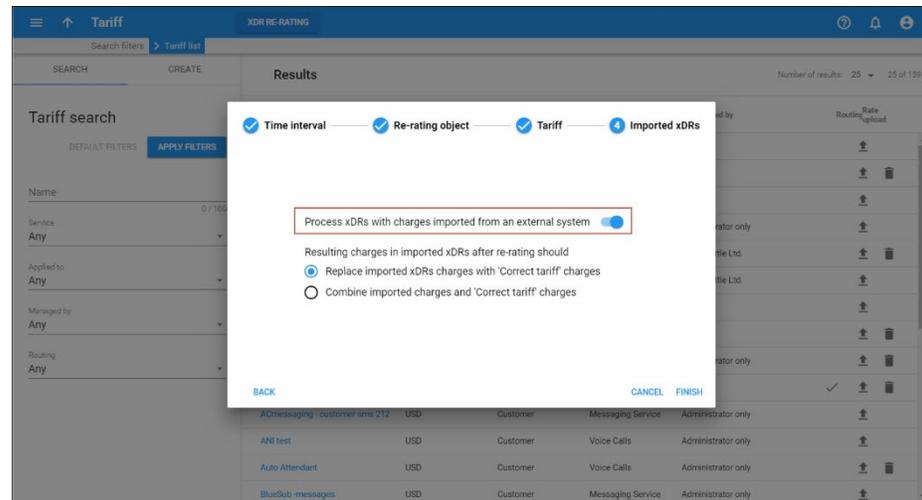
- If the rate exists both in the correct tariff and the override tariff – the override tariff rate applies;
- If the rate exists only in the correct tariff and is absent from the override tariff – the correct tariff rate applies.

An xDR re-rating using override tariffs can only be done for customers and resellers.

Re-rating of xDRs with external charges

If you import xDRs that already contain charged amounts to PortaBilling®, choose how PortaBilling® will handle those xDRs during the xDR re-rating procedure.

To re-rate xDRs that contain already charged amounts, enable the **Process xDRs with charges imported from external systems** option.



Then, during the xDR re-rating, PortaBilling® will either add the price from the tariff to the externally charged amount or reset the charged amount and calculate charges from scratch using the price from the tariff.

Call routing

When an entity on your network (for instance, a gateway or a SIP server) has to establish an outgoing call, it must find out where the call is being sent to. To do this, it can use its internal configuration (for example, dial-peer on the gateway) or some external entity (e.g. gatekeeper), or it may ask billing for a list of possible routes. This last method has several advantages: the routing configuration is in one central location, and billing can use information about termination costs to choose the best route (least-cost routing).

Thus, for SIP calls, PortaSIP® obtains the routing information directly from billing.

Control of routing on the Cisco gatekeeper (via GKTMP) or other equipment is not currently supported, but you can develop your own routing interface for PortaBilling®, which will communicate with PortaBilling® using RADIUS and then transfer this information to the originator in the correct format.

High-level routing overview

The following happens after the customer making the call has already been identified, the phone number has been translated into E.164, and all the required authorization checks have been performed:

- If the call is made to one of the access numbers for IVR applications, route the call to the PortaSIP® Media server and launch the IVR application.

- If the call is for one of the phone numbers provisioned in the system (there is an account with an account ID identical to the dialed number), direct the call to the PortaSIP® server where the SIP phone is currently registered. Also, provide a list of routes for follow-me numbers or a voice mailbox, so that the call can be redirected there if the SIP phone is not currently available (off-line, not answering, etc.).
- Otherwise, route the call to one of the gateways for termination, according to the routing rules specified in PortaBilling®.

Routing of SIP on-net (SIP-to-SIP) calls

As the first step in routing calculations, PortaBilling® checks if there is an account with an account ID equal to the dialed number (given in E.164). If such an account is found, it is clear that the call is designated for one of the subscribers provisioned on the system. PortaBilling® then extracts the latest known location for this account from the location database, in order to determine which PortaSIP® server the account is currently registered on, and routes the call there. The SIP server maintains information about all currently registered SIP user agents, so it is able to properly communicate with the SIP user agent handling this number.

Routing of off-net calls

You can have different vendors for terminating off-net calls. For example, you can terminate calls to the US either to AT&T, via a T1 connected to your gateway in New York, or to a remote gateway from Qwest.

Rate routing parameters

Ordinarily, tariffs define the termination costs for each connection to a vendor. If you create a tariff with the **Routing** type, a few more fields will be added to rates in that tariff:

- **Route category** – You can split rates into categories such as “Premium,” “Cheap,” etc. and use these categories in routing plans.
- **Preference** – Routing priority (0–10), higher values mean higher priority, 0 means do not use this rate for routing at all.
- **Huntstop** – Signals that no routes with a lower route category or preference should be considered.

This allows you to easily manage both termination costs and routing from a single location on the web interface. Thus, when such a routing tariff is associated with a connection, you can send calls for termination to all prefixes for which rates exist in the tariff.

Multiple routes

It is dangerous to have only one termination partner: if it is down, your customers will not be able to make any calls. Normally, you will try to find several vendors and enter their rates into the system. Each connection to a vendor (with routing tariff) will produce one possible route, and PortaBilling® will arrange these according to cost or your other preferences, thus providing fail-over routing.

Least cost routing

Least cost routing (LCR) is the method for computing routing lists based solely on vendors' cost. Thus, if you decide not to use any routing preferences (all vendors have the same preference), when a user makes a call, the LCR routing list will have the routes ordered by price from lowest to highest.

Compare the two illustrations below. The first one shows simple least-cost routing: all of the available carriers are arranged according to cost. In the second illustration, PortaSwitch® routing preferences are used, so that the administrator can re-arrange the routing. In this case, carrier D has been moved down the routing list, while carrier A has been moved up to the top of the list, thus becoming the first route.

Routing simulation

TEST

Dialed phone number *
8610234567

Routing plan
LCR (no routing plan assigned)

Node requesting routing
Any

Simulate with rates effective

Now

At the timepoint:
2020-01-03 17:04:00
Jan. 03 2020 17:04:00

#	CLD sent to vendor	Vendor	Connection IP/Node	Country Description	Tariff Destination	Price, USD	Route category	Preference	Huntstop
1	8610234567	Vendor D	Termination to... 167.22.45.67	China Proper	Vendor D 86	0.025	Default	5	
2	8610234567	Vendor C	Termination to... 125.55.75.32	China Proper	Vendor C 86	0.03	Default	5	
3	8610234567	Vendor A	Termination to... 153.64.75.86	China Beijing	Vendor A 8610	0.04	Default	5	
4	8610234567	Vendor B	Termination to... 123.34.45.56	China Proper	Vendor B 86	0.06	Default	5	
5	8610234567	Vendor F	Termination to... 188.45.73.42	China Beijing	Vendor F 8610	0.09	Default	5	
6	8610234567	Vendor E	Termination to... 155.27.32.12	China Proper	Vendor E 86	0.11	Default	5	

Routing simulation

TEST

Dialed phone number *
8610234567

Routing plan
LCR (no routing plan assigned)

Node requesting routing
Any

Simulate with rates effective

Now

At the timepoint:
2020-01-03 17:23:00
Jan. 03 2020 17:23:00

#	CLD sent to vendor	Vendor	Connection IP/Node	Country Description	Tariff Destination	Price, USD	Route category	Preference	Huntstop
1	8610234567	Vendor A	Termination to... 153.64.75.86	China Beijing	Vendor A 8610	0.04	Default	6	
2	8610234567	Vendor C	Termination to... 125.55.75.32	China Proper	Vendor C 86	0.03	Default	5	
3	8610234567	Vendor B	Termination to... 123.34.45.56	China Proper	Vendor B 86	0.06	Default	5	
4	8610234567	Vendor F	Termination to... 188.45.73.42	China Beijing	Vendor F 8610	0.09	Default	5	
5	8610234567	Vendor E	Termination to... 155.27.32.12	China Proper	Vendor E 86	0.11	Default	5	
6	8610234567	Vendor D	Termination to... 167.22.45.67	China Proper	Vendor D 86	0.025	Default	2	

Routing plans

Routing preferences in the rate allow you to specify that, for example, you would rather send a call to MCI than to T-Systems. However, this decision is “global”, and so will apply to all calls made in your system. But what if you would like to use MCI first for customer A, while T-Systems should be the first route for customer B, and customer C should be routed to MCI only?

This can be accomplished using **routing plans**. A routing plan is a combination of route categories in a specific order. It defines which categories of vendors are available for termination as well as in which order they should be arranged (route categories with a higher order value are tried first). For instance, in the example above MCI may be assigned as the “Normal” route category and T-Systems as the “Premium” category. After that, three routing plans will be created:

- **Quality** – Includes first “Premium” and then “Normal” route categories.
- **Ordinary** – Includes first “Normal” and then “Premium” route categories.
- **Cost-efficient** – Includes only “Normal” route category.

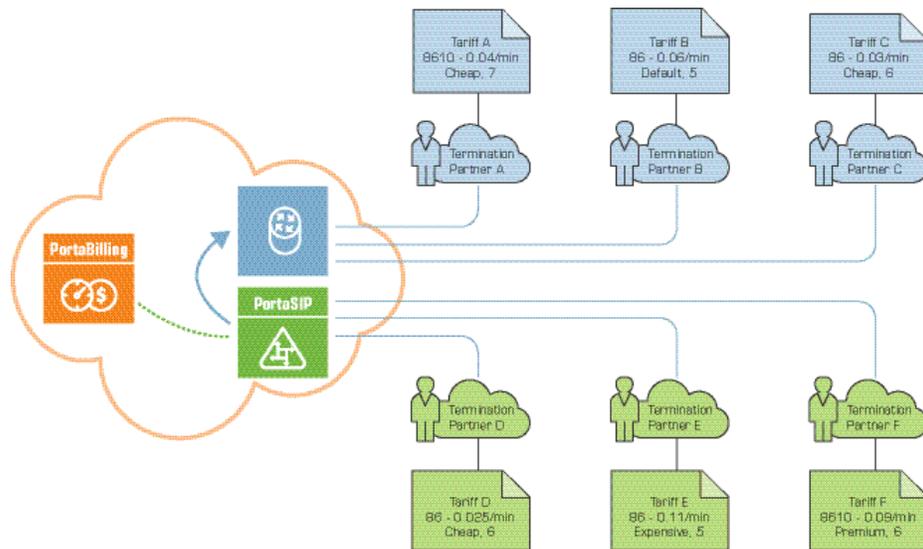
Thus, the system will offer a different set of routes depending on which routing plan is assigned to the account making a phone call. One routing plan, called **All Available Routes**, is always in the system, and is assigned to the account unless otherwise specified. This routing plan uses all available routes, regardless of their route category.

Route sorting

How exactly does PortaBilling® arrange multiple available routes?

1. By route category. Only route categories which are included in the routing plan will be used, following the order given in the routing plan.
2. If you have multiple route categories within the routing plan, you can either merge them into the same group by assigning them the same order value, or keep each one separate, with its own order value. Then routes within the same order group of route categories will be arranged according to preference.
3. For routes with the same preference, the system can arrange them according to cost (a comparison is made on the **Price_Next** rate parameter) so that cheaper routes will be among the first ones, or in random fashion.

Routing configuration example



Consider the following example. If you have:

- A “Standard” routing plan, which includes three route categories: “Default” (order 70), “Cheap” (order 40) and “Expensive” (order 10).
- Six vendors (A, B, C, D, E, F) with the following rates (prefix, route category, preference, price):
 - A. 8610, Cheap, 7, 0.04
 - B. 86 , Default, 5, 0.06
 - C. 86 , Cheap, 6, 0.03
 - D. 86 , Cheap, 6, 0.025
 - E. 86 , Expensive, 5, 0.11
 - F. 8610 , Premium, 5, 0.09

then, when a customer with this routing plan makes a call to **8610234567**, the system will arrange the possible routes as follows:

Vendor	Parameters	Comment
B	Default, 5, 0.06	The “Default” route category is first in the routing plan.
A	Cheap, 7, 0.04	This vendor has the highest preference in the “Cheap” category.
D	Cheap, 6, 0.025	This vendor has the same preference as vendor C, but a cheaper per-minute rate.
C	Cheap, 6, 0.03	
E	Expensive, 5, 0.11	This is the only vendor in the last route category.

(Vendor F was not included in the routing, since his route category is not in the customer’s routing plan).

Routing plans with profit guarantee

PortaBilling® supports a “profit-guarantee” mode for routing. If this mode is activated in a routing plan, all routes where the price per minute which your vendors apply to you is higher than that which you apply to your customers will be removed from the route list by the billing engine during real-time route calculation. Or, to put it in simpler terms, a route will not be selected if it is more expensive than what you charge your customers.

Routing plan selection

There are two possibilities for choosing which routing plan is used to find available vendors and terminate a call:

- Each account is assigned an individual routing plan. When the end user simply dials a number in the normal way, this routing plan will be used.
- If the end user wants to choose a specific routing plan for a particular call, he dials a routing plan selection code before the number.

For example, if he wants to use the “Premium” routing plan to call US phone number 12065551234, he would dial 23*12065551234, where 23* is the selection code defined as a property of the “Premium” routing plan.

Selection codes

The selection code is a property of a routing plan, and so obviously each routing plan has its own unique selection code. If you want to have a routing plan which cannot be specifically selected by an end user, simply leave the selection code field empty.

Order	Route category	Description
81	Economy	Low price, low quality
65	Cheap	

Dialing rules

By including a routing plan in the rating table definition in the product, you potentially allow every account with this product assigned to use this routing plan. If you want to prohibit a specific customer from selecting this routing plan for a call, you can disable it in the customer’s dialing

rules. In that case, calls will be routed using the account's default routing plan.

The screenshot shows the 'Dialing rule' configuration page. The left sidebar contains a menu with 'International', 'National', 'Exceptions', and 'Selection codes'. The 'Selection codes' section is expanded, showing 'Enable routing plan selection' (checked) and 'Remove selection code from CLD' (checked). The main content area is split into two panels: 'Dialing rule parameters' and 'Check for yourself'. The 'Check for yourself' panel displays a table of routing rules:

Local call	Domestic long distance call	International call
9856418 → 9856418	11 9856418 → 11 9856418	352 1 9856418 → 352 1 9856418

Required support outside PortaBilling®

In order to be able to use dynamic routing plan selection and have these calls charged properly, adequate support is required on the switch side, i.e. the node which processes these calls. This derives from the fact that, although the customer initially dials the selection code, this information is not included in the phone number sent to the vendor, and therefore is not visible in the call accounting information. Therefore, the node must provide this information to PortaBilling® as a separate attribute, so that the call can be charged properly.

PortaSIP® supports this functionality by default, while for switching equipment supplied by a third-party vendor it is recommended to consult with the vendor directly.

Routing algorithm

Based on the information above, the routing algorithm is as follows:

- The SIP server (or some other entity) asks PortaBilling® for a list of routes for a given number.
- PortaBilling® checks every tariff with routing extensions associated with a vendor connection to find rates matching this phone number. The best-matching rate is chosen in each tariff. This rate defines the routing parameters, i.e. it supplies a list of all vendors who can potentially terminate this call, including their cost and other related parameters.
- A list of possible termination addresses will be produced (this will include remote IP addresses for VoIP connections and IP addresses of your own nodes with telephony connections).
- This list will be sorted according to route category, routing preference and cost. Entries assigned a route category not

included in the routing plan and entries which do not satisfy the profit-guarantee criteria will be removed. If some entry has a huntstop defined, entries that have a lower preference or route category will be ignored.

- If there are several routes with an identical cost / preference, load sharing will be applied so that each potential route has an equal chance of being the first. Consequently, if your terminating carrier provides you with three gateways to send calls to, each gateway will, in the final analysis, receive approximately the same number of calls (33%).
- If adaptive routing is used, then routes to vendors currently penalized for not meeting the quality criteria will be moved to the very last positions in the routing list.
- A list of these IP addresses (with optional login and password for SIP authentication) will be returned to the SIP server. (To avoid extremely long call setup time, only a certain number of routes from the beginning of the list are returned; the default is 15, but this can be changed in the configuration).

Load balancing

The routing list, computed according to the costs and preferences of individual vendors, is “static” – meaning that each vendor will always maintain his respective position in the list (first route, second route, last route, etc.) *unless* the vendor changes his cost or the administrator manually adjusts the preference for that vendor. Let’s look at the following example:

Vendor	Preference	Cost (\$/minute)
A	5	0.05
B	5	0.06
C	5	0.15

In this case, after the sorting is done according to cost, Vendor A will occupy the first position in the list – and the first routing attempt will always be made to that vendor. Only if the call cannot be connected to Vendor A will it then be routed to Vendor B. This basically means that under normal circumstances, Vendor A will get most of the traffic, and a call will be routed to Vendor B only if all available capacity of Vendor A has been exceeded or if there is some temporary problem with connecting the call. Although this scheme offers maximum profit, conducting business this way is usually undesirable in real life. If the cost of several routes is comparable, it is better to allow a portion of the traffic to go further down the routing list via other carriers (Vendor B in our example). This will ensure that the other routes are maintained in an active state, that there is a sufficient amount of data for producing quality reports and

that your contracts with those vendors will not be canceled due to inactivity.

The idea is to allow a portion of all traffic to be directed to the vendors who are further down the routing list – though of course, this can only be done if the cost difference between the vendors is acceptable. In our example, Vendor C has a price per minute of \$0.15 – so this route should obviously not be mixed with those of Vendors A and B, since the difference in cost (\$0.10/min) is unacceptable.

By setting up the **load balancing threshold** in the routing plan, the administrator specifies the maximum acceptable difference in cost between two carriers, thereby indicating when routes to these carriers can potentially start “swapping” positions in the routing list.

How does this value affect the route sorting in exact terms? After the list of potential routes is produced, but before the actual route re-arrangement according to cost takes place, each route cost is slightly changed by a random value. The adjusted cost may be increased or decreased compared to the original one, but the spread between the minimum possible value and the maximum possible value will be specified as the **load balancing threshold**.

For instance, when the load balancing threshold is set to \$0.02, then Vendor A’s cost (\$0.05) may be increased or decreased by \$0.01 (one half of \$0.02), changing the cost to a random value between \$0.04 and \$0.06. Naturally this adjusted cost is purely fictional; these changes would only take effect during the sorting of routes and would not affect the actual call billing.

As a result of these cost fluctuations, when the list of routes is sorted according to cost, the route sequence “shuffles” and the probability of two routes swapping places increases when they have similar prices. If in our example, above, we set the load balancing threshold to \$0.02, the adjusted cost of A would be in the range of \$0.04 and \$0.06, and the adjusted cost for B would be in the range of \$0.05 and \$0.07. If for instance, the new cost of A happens to be \$0.053 and the adjusted cost for B happens to be \$0.052 – Vendor B would move into the first route position.

Overall, in about 25% of all cases both values would be randomly distributed in the same range (\$0.05-\$0.06) and therefore, B would become the first route for about 12,5% ($25\% / 2$) of all call attempts. In comparison, if B’s cost were \$0.051 – this algorithm would result in B receiving an almost equal share (49.5%) of all calls that A is receiving, since adjusted cost values for A and B would be randomly distributed throughout two nearly 100% overlapping intervals (\$0.04-\$0.06 and

\$0.041-\$0.061) and thus each vendor would have nearly the same chance of his adjusted cost being the lowest.

This process of randomly adjusting the cost followed by re-arrangement is done for every new call. So for each call the actual routing list may look slightly different from previous lists and the overall call volume is properly distributed among multiple vendors.



Note that load-balancing only applies to routes with the same preference and route category that have been assigned the same order in the routing plan. This is to prevent load-balancing (when undesirable) by simply increasing or decreasing the preference of a specific route. So if in our example the route to Vendor A has a preference of 6 and the route to Vendor B still has a preference of 5 – load-balancing would not happen, and A would remain in the first route position.

Routing and connection utilization

Aside from the price difference, another important consideration when arranging multiple routes is the number of already connected calls for a particular connection and the maximum number of calls it can potentially handle. Without this, a connection to a carrier may become overloaded, resulting in decreased sound quality and a higher number of failed call attempts, etc., while other carriers with just a slightly higher price may be underutilized.

This is why PortaBilling® offers you the ability to factor in connection utilization when computing the routing list and perform load balancing among multiple connections accordingly. Load-balancing based on connection utilization may change route precedence even in cases where the price difference between vendors is too high for normal load-balancing (as described earlier).

For each connection, the administrator defines two parameters: the **maximum connection capacity** and the **load threshold** (on PortaBilling® web interface it is the **Start Utilization Balancing After** attribute).

Connection SAVE CHANGE STATUS

Termination to Z-Telecom > **General configuration**

✓ **Termination to Z-Telecom** 24 / 32

Service type: Voice

Type of connections: SIP

Direction: To vendor

Tariff: Termination to Z-Telecom

Routing criteria: None

Internal:

General configuration

Identify gateway by: IP

Remote IP*: 192.28.187.205 : Port 14 / 19

Gateway ID:

Service policy: Default

Capacity*: 60

Vendor authorization: None

General configuration >

Service configuration >

Load graphs >

Connection SAVE CHANGE STATUS

Termination to Z-Telecom > **Service configuration**

✓ **Termination to Z-Telecom** 24 / 32

Service type: Voice

Type of connections: SIP

Direction: To vendor

Tariff: Termination to Z-Telecom

Routing criteria: None

Internal:

Service configuration

Protocol: UDP

Use for:

Limit utilization by capacity:

Start utilization balancing after: 50 2 / 255

RTP proxying: Optimal

Caller identity: Do not supply

Translate CLD: Do not translate, send E.164 to the...

Tech prefix:

General configuration >

Service configuration >

Load graphs >

As long as the number of connected calls remains below the load threshold, the connection is considered to be operating in normal mode (sufficient available capacity), and calls are routed to this connection based solely on cost and assigned preferences (if the load-balancing threshold is defined, the shuffling of routes with similar prices may also take place).

If the number of simultaneous calls established via the connection exceeds the load threshold, the connection is considered to have surpassed the desired limit and therefore some of the load shifted to other carriers that possess an acceptable price difference. This is done using an algorithm similar to the one used for load-balancing (described in the previous chapter). The administrator sets the **overload handicap** parameter (defined per routing plan) that specifies how far above the desired limit a load may “weigh down” the connection. A portion of this value (proportional to the actual load on the connection) is added to the route cost during route sorting so that the vendor’s cost (whose route was the first ‘unconditional’ route due to its lower price) now becomes comparable with the other vendors’ costs, and load-balancing can commence.

If the connection load increases even further, the added “handicap” value increases as well. The calculated vendor’s cost (that is used during sorting) would rise higher and higher, leaving less chance for this vendor to be the first route.

Finally, when the quantity of calls in progress reaches the maximum connection capacity threshold, this connection becomes excluded from the routing process; no further calls will be routed there until some already connected calls are disconnected.

Let’s extend the example used for load-balancing to consider connection utilization. There are three vendors:

Vendor	Cost	Maximum Capacity	Load Threshold
A	\$0.05	100	75
B	\$0.07	250	150
C	\$0.10	150	100

The routing plan setup is as follows:

Parameter	Value
Load-balancing threshold	0.01
Overload handicap	0.05

Initially, when the load across all connections is minimal, the routing order is determined only by the vendor’s cost. Since the difference in cost is more than the allowed load-balancing threshold (\$0.01), the routing order will always be Vendor A, Vendor B and then Vendor C.

How does the situation change when some connections surpass the desired load limit?

Vendor	Cost	Max. Capacity	Load Threshold	Current Load	Handicap	Adjusted Cost
A	0.05	100	75	83	0.016	0.066
B	0.07	250	150	30	0	0.07
C	0.10	150	100	5	0	0.10

As you can see, Vendor A has exceeded the desired load capacity (75, specified by the load threshold) by 32%. This is calculated as the ratio of total number of calls above the desired capacity against the difference between desired and maximum capacity, in this case $(83 - 75) / (100 - 75)$. So 32% of the overload handicap ($\$0.05 * 32\% = \0.016) is added to Vendor A's cost and the adjusted cost becomes \$0.066. Since the difference between Vendor A's and B's costs is less than the load balancing threshold, load-balancing will commence and some percentage of calls will be routed to Vendor B. Let's continue the example and assume that Vendor A's load has increased even further:

Vendor	Cost	Max. Capacity	Load Threshold	Current Load	Handicap	Adjusted Cost
A	0.05	100	75	90	0.03	0.08
B	0.07	250	150	155	0.0025	0.0725
C	0.10	150	100	10	0	0.10

Since Vendor A's load is approaching the maximum allowed limit, they are penalized even further – now with an adjusted cost of \$0.08. Vendor B's cost is also starting to be adjusted. Since it is only 5 calls above the load threshold, it is quite low compared to the maximum potential number of calls above the load threshold, which is $250 - 150 = 100$ calls. So the overutilization is $5 / 100 = 5\%$, and thus the handicap effect is minimal ($\$0.05 * 5\% = \0.0025). Now Vendor B will actually become the first route for the majority of calls, since its adjusted cost (\$0.0725) is lower than Vendor A's (\$0.08).

Vendor	Cost	Max. Capacity	Load Threshold	Current Load	Handicap	Adjusted Cost
A	0.05	100	75	98	0.046	0.096
B	0.07	250	150	200	0.025	0.0925
C	0.10	150	100	20	0	0.10

When the load on both Vendors A and B increases, their adjusted costs increase to the point that Vendor C also begins to participate in the load-balancing.

Vendor	Cost	Max.	Load	Current	Handicap	Adjusted
--------	------	------	------	---------	----------	----------

		Capacity	Threshold	Load		Cost
A	0.05	100	75	100	0.046	N/A
B	0.07	250	150	230	0.04	0.11
C	0.10	150	100	20	0	0.10

And finally, when the load on Vendor A's connection reaches the maximum allowed limit, this connection will be excluded from call routing. All further call attempts will be routed to Vendors C and B (with C becoming the first route because of its lower adjusted cost).

Routing override

Sometimes it is necessary to make minor adjustments to one or several routes. This configuration can be done with the help of the Routing Override functionality for individual routing plans.

For example, a pool of connections is created in which 70% of the calls should go to carrier A, 20% to carrier B and 10% to carrier C. Suppose that the system has a total of 1,000,000 minutes. The administrator expects that 700,000 minutes will go to carrier A, 200,000 minutes to B and 100,000 minutes to C.

In order to achieve this, the system will randomly change the routing list for each call and place one of them in first place according to the specified percentages. If you look at the statistics at the end of the day, you see that you obtained the required ratio.

Once you have determined which route will be in first place, you will decide what to do with the other routes. Therefore, there are two cases of routing override that are configured using the pool:

- All routes are on the routing list, and all will be used for calls, thereby reducing the number of failed calls.
- Only the first route is used for all calls, the other routes are ignored.

Several types are used for configuring the pool of connections:

- Consistent routing to one carrier (e.g. traffic to Vietnam-Mobile from customer A always goes to carrier XYZ).
- Sequential routing – The administrator creates a list of connections for a destination group in a desired routing order, and calls are routed according to this sequence, before (or instead of) applying “normal” LCR routing. An entry in the override list can be a percentage-share connection pool, where each connection has a chance to be the first route, proportionate to the assigned percentage value (e.g. 25 percent of cases send traffic via the connection Europe Premium, 40 percent of cases send via the

connection Termination to UK, and the rest (35 percent), are sent based on LCR).

The configuration examples for the above-mentioned types of routing override can be found in the [Unified PortaSwitch Handbook Collection](#).

Routing margin tariff

In some cases, setting a single profit margin for all destinations is not flexible enough for obtaining the best results. For example, the service can be profitable when a service provider sells it to customers at “10% above the carrier’s price” but market research shows that certain destinations can be charged a higher rate and customers will buy the service anyway. The service provider sets a higher price for such destinations (e.g. 25% above the standard carrier’s price), and now wants to exclude those carriers that produce less profit from the routing.

PortaBilling® supports routing margin tariffs that offer service providers an opportunity to maximize their profits. A routing margin tariff is a special tariff to be used instead of a real customer tariff exclusively in order to compare carriers’ prices against it. Only carriers having prices lower than or equal to those specified in this tariff will be selected for routing.

Consider the following example:

An administrator is configuring the profit-guarantee settings and sets the minimum profit to 10%, specifying that it must be calculated based on the routing margin tariff.

The user calls a number in the UK, 44125551155. According to the user’s product, the regular rate for destination 4412 is \$0.40. If there were no routing margin tariff defined, PortaBilling® would use those carriers whose price for this destination is either lower than or equal to $\$0.40 - 10\% = \0.36 .

However, the routing margin tariff is set so PortaBilling® first checks whether there is a rate for 4412. PortaBilling® finds the rate for this destination defined as \$0.10 so only carriers having a price for 4412 of $\$0.10 - 10\% = \0.09 and lower will participate in the routing.

A regular customer tariff is used then solely to charge the customer. However, if a certain destination is not present or is prohibited in the routing margin tariff, the price from the regular customer tariff is used for calculating the profit-guarantee outcome.

The screenshot shows a web interface for 'Routing simulation'. On the left, there are input fields for 'Dialled phone number' (8610234567), 'Routing plan' (Standard), and 'Node requesting routing' (Any). Below these are options to 'Simulate with rates effective' (Now). The main area is a table titled 'Routing list' with 6 rows of results. The table columns are: #, CLD sent to vendor, Vendor, Connection IP/Node, Country Description, Tariff Destination, Price, USD, Route category, Preference, and Huntstop.

#	CLD sent to vendor	Vendor	Connection IP/Node	Country Description	Tariff Destination	Price, USD	Route category	Preference	Huntstop
1	8610234567	Vendor B	Termination to... 123.34.45.56	China Proper	Vendor B 86	0.06	Default	5	
2	8610234567	Vendor F	Termination to... 188.45.73.42	China Beijing	Vendor F 8610	0.09	Default	5	
3	8610234567	Vendor A	Termination to... 153.64.75.86	China Beijing	Vendor A 8610	0.04	Cheap	7	
4	8610234567	Vendor D	Termination to... 167.22.45.67	China Proper	Vendor D 86	0.025	Cheap	6	
5	8610234567	Vendor C	Termination to... 125.55.75.32	China Proper	Vendor C 86	0.03	Cheap	6	
6	8610234567	Vendor E	Termination to... 155.27.32.12	China Proper	Vendor E 86	0.11	Expensive	5	

Adaptive routing

PortaBilling® allows you to automate the process for controlling the call quality which becomes increasingly important today. That is, if you want to evaluate acceptable vendors for terminating VoIP calls, there is no need to hire numerous human operators or network engineers, who will track and analyze the specific route. All you need is to implement adaptive routing model.

The main idea of adaptive routing is to dynamically measure a vendor's quality parameters, and adjust the routing priority accordingly. These quality requirements are predefined in the form of threshold parameters on the **Routing criteria** page, and are then automatically applied to specific vendors. Any vendor who fails to satisfy your quality requirements will go to the “penalty box” – the very bottom of the routing list. This means that the system will try first to terminate calls using other carriers (with a good quality evaluation). However, if all of them fail or are unavailable, the “penalized” carrier will have a chance to terminate the call. It is still better to send a call via an inferior vendor than to have it fail completely.

Thus, if the quality requirements are applied, a carrier's place on the routing list is determined not only by the route category, the assigned preference value, and the cost parameters (LCR model), but additionally by quality criteria.

For effective quality measurement, the following key control parameters are used:

- ASR (Average Success Rate) – The number of successfully connected calls divided by the total number of call attempts.
- PDD (Post Dial Delay) – The time interval between the connection request to a vendor and ring-back.
- ALOC (Average Length of Call).
- PPM (Profit per Minute) – The aggregated profit, i.e. the difference between the actual charged amounts in the CDRs for your customers and vendors.

Each of the above parameters requires two values, which define the warning and penalty thresholds, respectively. When the value of a parameter reaches the predetermined threshold, the administrator receives an email alert about the latest connection threats. Moreover, the administrator can track the current connection status on the **Tracking** page. This status is represented by different colors, as follows:

-  GREY – The number of calls is not enough to apply filtering differentiation.
-  GREEN – The route meets the quality requirements.
-  YELLOW – The route is active, but some of its quality parameters are outside the warning thresholds.
-  BLOCKED – This route is currently being penalized.

NOTE: The penalized route will stay in the “penalty box” for the period of time specified in **Penalty Time**, and then will be unblocked automatically. Alternatively, you can unblock the penalized route manually by clicking the **Unblock Now** button.

-  RED – The route was manually unblocked; this status will remain until the next time interval for computing statistics.

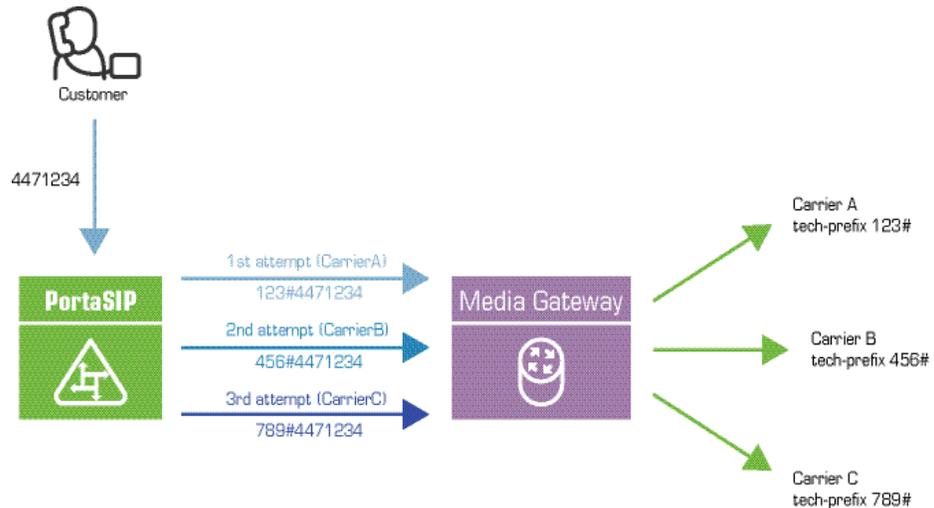
Routing to an external media gateway

A common situation is when several different carriers are connected via PRI trunks (or SS7 interface) to the same media gateway, which does the conversion from VoIP to PSTN. When a call is sent to the gateway, the question is which particular carrier should be used to connect this particular phone call (or if multiple carriers are used – in which order should they be tried)?

Most of the gateways support configuration of routing for outgoing calls via dial-peers or similar method, where the outgoing carrier or trunk group is chosen based on matching the dialed number (e.g. 4471555123456) to a set of prefixes (e.g. 447 or 4471) assigned to outgoing trunks. So it is possible to maintain the routing configuration on the gateway, but it is not taking into consideration the costs of sending calls to a particular vendor, and any change in the routing requires manual re-programming of the routing table on the gateway. Needless to say, this method significantly increases the amount of work for the administrator, reduces routing flexibility and is very likely to generate routing errors, which could seriously damage revenue of the service provider, since calls are not being sent to the carrier with lowest costs.

So another approach is recommended: the media gateway is configured with the bare minimum of the routing information. Each outgoing carrier (or trunk group) is assigned its own prefix (e.g. 123# for carrier A and 456# for carrier B). When PortaSIP® sends a call to media gateway, it

appends the specific carrier's prefix to the dialed number, so only that carrier is used to transport this call (as illustrated below).



As a result, the call routing is controlled at the single location, and one has to maintain a single set of carrier rates and related configuration.

Routing in a multi-node scenario

Sometimes a network consists of multiple nodes (for instance, if you have two PortaSIP® servers installed), each of them capable of doing call routing under the control of PortaBilling®. In this case, two scenarios are possible:

- The node which receives the original call from the customer proceeds with routing of the call and delivers it to the “destination” carrier. This is the simplest scenario, and usually it provides the best results in terms of your network’s usage (since the call always travels the shortest possible path, with a minimum number of hops on the network). In this case, though, each node communicates with the carrier directly, which may require additional maintenance activities on your side (e.g. if the carrier uses authorization by IP address, you will need to provide him with a list of all your nodes’ IP addresses and keep that list updated as your network changes).
- The node which receives the original call from the customer does the authorization (to verify the identity of the caller) and sends the call over to the “master” node. The master node then obtains the actual list of routes for this call and establishes the connection to the terminating carrier.

The last scenario increases traffic within your network, but it allows you to handle advanced scenarios – for instance, when you have multiple PortaSIP® nodes deployed in different countries (and their IP addresses

change often), but only your central PortaSIP® node, located in your main collocation facility, is able to communicate with the carrier because of IP address restrictions.

Routing to carriers, which only support H323

If the number of H323-only carriers is not too high, it would make sense to configure the system as follows:

You can use Cisco IPIP GW, or any of the SBCs available on the market as a SIP-H323 converter. The idea is to configure it in such a way that any incoming SIP call where the destination number starts with a certain tech-prefix will be sent out to a particular carrier. For example, if the destination number in the incoming call starts with 1#, the call will go out to IP address 6.7.8.9 (SuperNet); if the destination number starts with 2#, to IP address 1.2.3.4 (X-Telecom); and so on. Although this converter is (strictly speaking) part of your network, there is no need in this case to define it as a node in PortaBilling®, since it will not be exchanging any information via RADIUS.

In PortaSwitch® you will configure your SIP carriers as usual. When you create a connection for an H323 carrier, you will specify the IP address of your SIP-H323 converter as the **Remote IP Address** and then assign a tech-prefix to this connection. In this way, although calls for multiple carriers will go from PortaSwitch® to the same IP address, the system can still recognize which carrier is used in each case.

Let's consider the following example for a better illustration of how this process works. Assume you have 4 carriers, all capable of terminating calls to the UK. Their respective rates per minute are: carrier A – 0.02, carrier B – 0.03, carrier C – 0.05 and carrier D – 0.10. Carriers B and D are H323 only, so for them connections are configured on the SIP-H323 converter (with IP address 192.168.0.3), with tech-prefix 2# and 4#, respectively. A customer dials phone number 4412345678. The call arrives to the PortaSIP® node, which, upon successful authorization, receives the following routing list:

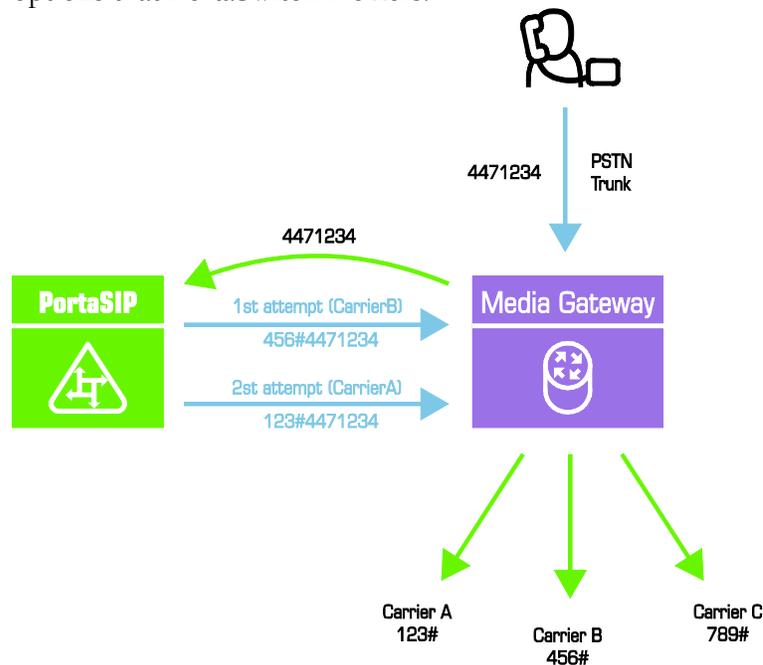
```
Route 1: 4412345678 @ <IP address of carrier A>
Route 2: 2#4412345678 @ 192.168.0.3
Route 3: 4412345678 @ <IP address of carrier C>
Route 4: 4#4412345678 @ 192.168.0.3
```

PortaSIP® will first try to establish a connection via the first route and, if that fails, it will send the call via the second route, to the SIP-H323 converter, which in turn will attempt to establish the call with carrier B. If a connection is not possible, after receiving an error message from the converter PortaSIP® will try the next available route, and so on.

This method of configuration allows the SIP-H323 converter resource to be used only when it is actually required (since SBCs are usually licensed on a per-port basis), thus minimizing total network infrastructure costs.

Intelligent routing on a third-party media gateway / SBC

A common situation faced by service providers today is when a piece of communications equipment on the network (VoIP/TDM switch, media gateway, SBC, etc. – referred to below as **the switch**) is not capable of doing real-time call authorization or intelligent call routing like PortaSwitch®. How can you still perform call processing on such equipment in conjunction with the advanced call control and routing options that PortaSwitch® offers?



In this scenario, the most straightforward method would be the following:

- The switch is configured to accept an incoming call initiation request and send it to PortaSIP®.
- PortaSIP® receives an incoming call and sends a request to PortaBilling® to perform caller authorization and obtain the routing list.
- If authorization fails, the call disconnect command is sent back to the switch, and the incoming call on the switch is dropped.
- In case of successful authorization, PortaBilling® computes the routing list (using LCR, preferences, settings in the customer's individual routing plan, the current status of adaptive routing information for carriers, etc.). The resulting list of routes is delivered to PortaSIP®, which then starts processing them. If

there is a carrier directly connected to the switch, PortaSIP® sends the call back to the switch, adding a tech-prefix to the destination number to indicate the specific trunk group this call should be routed to. If the first route fails, another call is sent to the switch (indicating a different trunk group), and so on. When the call is established, PortaSIP® acts as a signaling-only node, so there is no effect on the codecs, sound quality, or other voice media attributes. Finally, when the maximum allowed time has been reached, PortaSIP® sends a call disconnect request to the switch.

This allows the use of call authorization, advanced real-time routing and control of the customer's credit limit, even with equipment which does not directly support these options.

Number translation

Your network numbering plan

Telecommunication service providers use the E.164 format in their networks (e.g. PSTN, mobile or SIP). This unifies the numbering format among different networks, simplifies routing and indicates that phone numbers belong to residents of a certain geographical area.

Phone numbers in E.164 format (e.g. 12391232967) start with a country calling code (e.g. 1), followed by an area code (e.g. 239) and a user's local number (e.g. 1232967).

It is common practice to write phone numbers in E.164 format (e.g. on a company's web site or a business card) with the '+' prefix (e.g. +12391232967). The '+' indicates an international calling code (e.g. 00 or 011), though customers usually dial phone numbers from any phone (e.g. mobile, landline or SIP) without the '+.'

When PortaBilling® receives the number with the "+" prefix (e.g. +1234567890,) it recognizes that the number following the prefix is in the E.164 format, so no local / custom dialing instructions apply.

This information gives administrators the ability to create internal dialing rules and adjust their customer's / vendor's numbering formats to the unified E.164 format used in PortaBilling®. For this purpose, they use **translation rules**.

Translation rules

Very often you need to use different numbering systems on your VoIP network. For example, vendor A requires you to send all phone numbers with the technical prefix **12345#**, Vendor B requires that all phone

numbers be in the format **011<country-code><area-code><number>**, and when you send calls to the local telecom (Vendor C) you have to strip the country code and area code off completely. Of course, you would like to have this information in a unified format in your billing, so that you do not need to enter multiple phone prefixes for the same destination. Another common situation is when customers located in different countries or regions (but all managed on the same PortaBilling® platform) dial phone numbers in different formats, since every customer dials a number in the way he or she is used to.

Thus the solution is to maintain all the numbers inside your network and PortaBilling® in a unified format – E.164 – and use **translation rules** to deal with situations when numbering formats differ.

There are two main classes of translation rules:

- Customer-specific *dialing rules*, used for converting numbers dialed by end users to the PortaSwitch® format – E.164 format;
- Vendor-specific *connection rules*, used for converting numbers from the E.164 format to one required by a particular vendor for further call processing.

A translation rule is actually a set of Perl regular expressions which are applied to a phone number. So if you are familiar with regular expressions, you can make all kinds of transformations with phone numbers; for more information, see the [PortaBilling® online help](#). If you are not familiar with regular expressions, you can use a dialing rule wizard which will construct the correct rule based on the parameters you provide, such as country or area code. This allows you to easily manage a network with many different numbering plans for your customers and vendors. Also, translation rules can be applied to ANI (CLI) numbers, allowing caller information to be sent to a vendor in the format he requires (e.g. 10-digit phone numbers for US callers).

This is how dialing and connection rules work in conjunction: When a customer makes a phone call, it arrives from outside your network and the customer's dialing rule is applied. Thus the dialed number is converted from the customer's format (e.g. 0114202123456) into the unified format (4202123456). Next, all the billing calculations (authorization, routing, etc.) are done, and then, if it appears that this call should go to a vendor who requires a special format, the outgoing connection rule associated with this vendor's connection is applied, and the number is translated into the required format (e.g. 123#4202123456). Finally, when the call is disconnected and PortaBilling® receives information about the completed call to phone number 123#4202123456, the vendor's connection rule is applied, so that the number becomes 4202123456. This is the number that will be stored in the xDR information.

Dialing rules

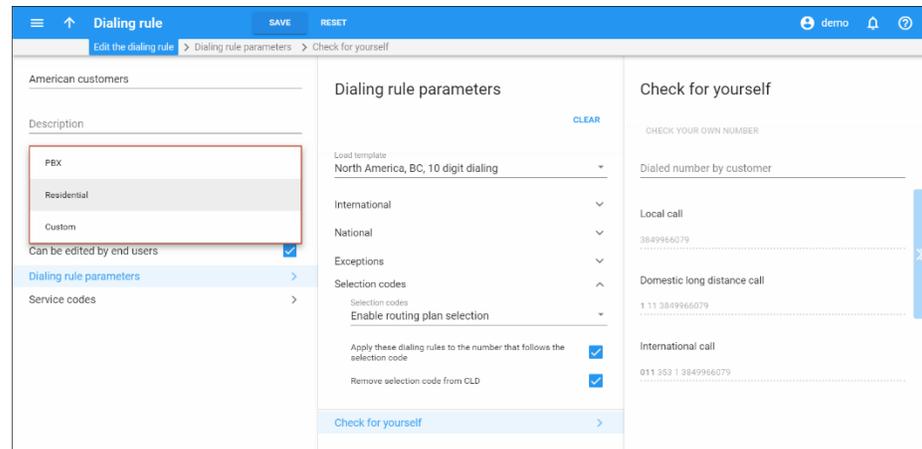
Frequently a customer will have his own numbering format, for example, dialing 00 for international numbers and just the phone number for local calls. Dialing rules allow you to convert a number from a format specific to a particular customer. Typically, dialing rules are either specified for the entire IP Centrex environment (provisioned as customer in PortaBilling®) or for an individual phone line (provisioned as account). Customers can manage their dialing format themselves via the self-care interface.

Dialing rules have two applications:

- When a number is submitted for authorization, these rules will be applied, with the resulting number used to search rates. Thus, if your customer dials 0042021234567, the billing engine will convert it to 42021234567 and find the correct rates for the 420 prefix.
- This number will be returned to the node which requested it.

To simplify the dialing rule configuration, they are divided into the following types:

- **Residential** – This is the simplest dialing rule type since it emulates the dialing habits of residential customers. In this dialing rule you define the international and national codes plus the special prefixes that residential customers must dial to access the network and make a call (e.g. when making calls from an IP Centrex environment).
- **PBX** – This is a more extensive dialing rule. It includes the dialing parameters of the residential type plus the service codes (e.g. *70 for call parking) used within the IP Centrex environment.
- **Custom** – This type allows you to define your own dialing rules by using Perl regular expressions.



The screenshot shows the 'Dialing rule' configuration page. The left sidebar has 'American customers' selected. The main content area is split into three columns. The first column shows a list of dialing rule subtypes: 'FBX', 'Residential', and 'Custom'. The 'Residential' subtype is highlighted. The second column, 'Dialing rule parameters', has a 'Load template' dropdown set to 'North America, BC, 10 digit dialing'. Below it are 'International' and 'National' dropdowns. There are also 'Exceptions' and 'Selection codes' sections. The 'Selection codes' section has 'Enable routing plan selection' checked. The third column, 'Check for yourself', has a 'CHECK YOUR OWN NUMBER' dropdown set to 'Dialing rule parameters'. Below it are 'Local call' and 'Domestic long distance call' sections, each with a number field. The 'Local call' field contains '38-49966079' and the 'Domestic long distance call' field contains '1 11 3849966079'. There is also an 'International call' section with the number '011 553 1 3849966079'.

Since dialing rules in North America are different from those in other countries and contain either 7- or 10-digit numbers, a separate **North America** dialing rule subtype exists that allows the administrator to easily configure and manage dialing rules for North American customers.

Connection rules

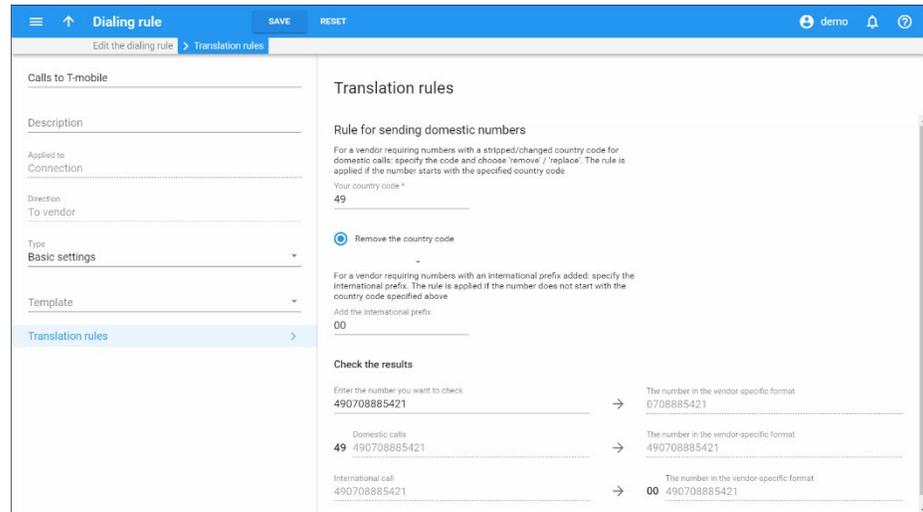
If your vendor requires a special number format (e.g. tech-prefix), it is possible to apply a connection rule to convert the number. When the billing engine returns a list of routes, the phone numbers for routes for this connection will be converted.

There are two types of connections:

- Outgoing connections via which calls leaving your network are sent to a vendor and
- Incoming connections via which calls to DID numbers are delivered from a vendor to your network.

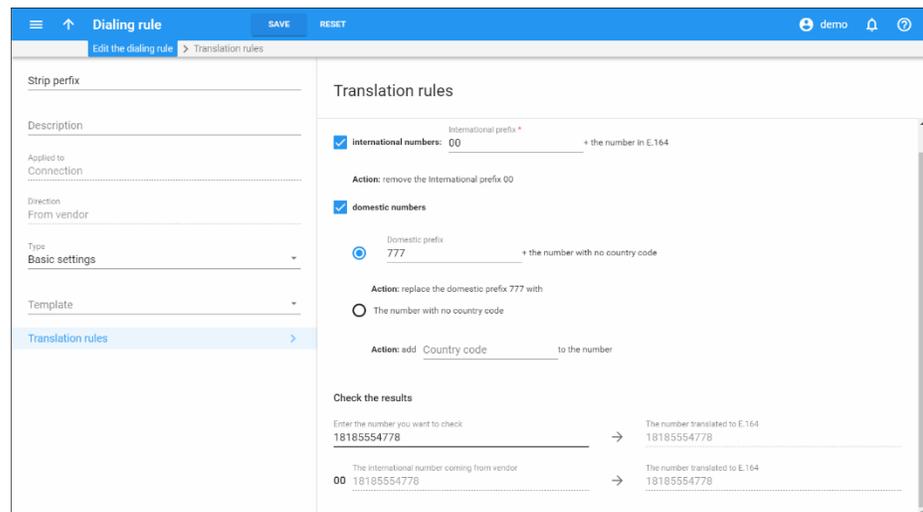
Connection rules differ depending on connection type and consequently, contain different sets of options. Therefore, when creating a connection rule, the administrator needs to specify only those options relevant for a particular connection type. This significantly simplifies the rule creation process and eliminates configuration errors.

Outgoing connection rules



When the call has been terminated to the vendor in a vendor-specific format, it will be reported to billing in this same format (e.g. 01142021234567). Now it is necessary to convert this number to the proper format for billing (4202134567), which may be done using connection rules. These rules will be applied to all calls which go through a given connection (even those routed there using dial-peers or other external tools).

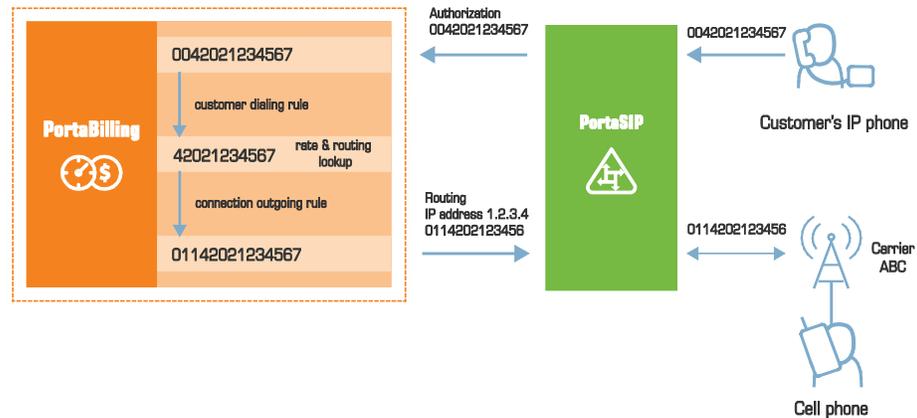
Incoming connection rules



Connection translation rules are also used on Calls from Vendor via SIP connections, when a number received from a vendor (e.g. 2065551234) must be translated into E.164 format, so that it will properly match an account whose ID is in E.164 format (12065551234).

Number translation on your network

Below is an illustration of how different translation rules are applied during a call.



1. The customer dials a phone number on his SIP phone. He enters the number in the same format he uses on his conventional phone, i.e. 0042021234567.
2. The number is delivered to the PortaSIP® server and translated using the **customer's dialing rule**, which states that the international dialing prefix for this customer is 00. So the number becomes 42021234567 (E.164 format). This number is used to search for the customer's rate for this destination.
3. PortaSIP® then requests routing for this number. Carrier ABC is defined for terminating calls to the Czech Republic in PortaBilling®. However, this carrier requires the number to be in US dialing format, so the international number must be prefixed by 011. An **outgoing translation rule** `s/^/011/;` to carrier ABC has been defined, and is now applied to the phone number, with the result 01142021234567. Note that there may be several carriers who can terminate this call, each with its own numbering format. In such a case there will be several alternative routes with different phone numbers.
4. PortaSIP® attempts to establish a connection to remote gateway 1.2.3.4 using phone number 01142021234567.
5. After the call is completed, PortaSIP® sends an accounting request to PortaBilling®, stating that the call to remote gateway 1.2.3.4 has just been completed. PortaBilling® finds a connection to vendor ABC with remote IP address 1.2.3.4, and applies the **translation rule** `s/^011//;` for this connection in order to convert the number from the vendor-specific format into your billing format. Thus 011 is removed from 01142021234567, and the number becomes 42021234567. PortaBilling® searches for the vendor and customer rates for this number and produces the CDRs.

Favorite number billing

Although it represents a totally different business case, this type of billing uses the same internal architecture of number translation plug-ins, described in the previous section. The only difference is that, in this case, the dialed number is checked against a list of “favorite” numbers defined for each account. If a match is found, the call is rated according to a special rate for the **FAV** destination, defined in the customer’s tariff. This allows you to offer customers a “call friends & family cheaper”-type service.

Premium number billing

In a premium number billing model, PortaBilling® investigates (either in a local database or an external one) to determine a number’s routing destination and which rating code must be used for charging your customer.

US inter / intra state billing

PortaBilling® offers a simple solution for identifying and billing for US inter / intra state VoIP traffic. In order to properly apply LCR and calculate call costs in accordance to rates applied by vendors, PortaBilling® determines whether the caller resides in the same state as the called party and then applies the correct rate.

When a user from the United States dials a US phone number, PortaBilling® first checks if both the number calling and the called number start with a “1.” If so, then PortaBilling® searches for both destinations in the database. If these destinations are in the same state, PortaBilling® applies charges according to the intra-state rate; if the destinations are in the different states, the inter-state rates are applied.

The following taxation plugins are supported in PortaBilling® to calculate customers’ taxes for the inter / intra-state calls:

- Avalara;
- EZtax (former BillSoft);
- SureTax.

For more information about taxation plugins, please refer to the [Processing taxes](#) section of this guide. Find guidelines how to configure inter/intra state billing in the [Billing for US Inter / Intra State Calls](#) handbook of the PortaSwitch Handbook collection.

DID provisioning and billing

DID inventory management

The DID inventory in PortaBilling® facilitates working with the customer's list of DID numbers, even if they are not currently assigned to any accounts. It also allows customers to manage their DID lists via the self-care interface.

Customers can see a list of currently provisioned numbers (and which extensions they are assigned to) on the “DIDs” tab of the customer self-care interface.

This approach can be used for geo and non-geo DIDs as well as for toll-free and premium numbers.

When working with DID inventory and provisioning DID numbers it is very important for service providers to have a clear idea of which costs they bear and which charges they apply to their customers.

DID number cost

DID providers traditionally offer DIDs on a “fixed monthly fee” basis, with no per-minute charges. The providers include access for obtaining “buying” rate information for a specific number via XML API (e.g. phone number 12065551234 will cost \$5/month). This is the rate the DID provider will charge the ITSP that allocates the number to its network.

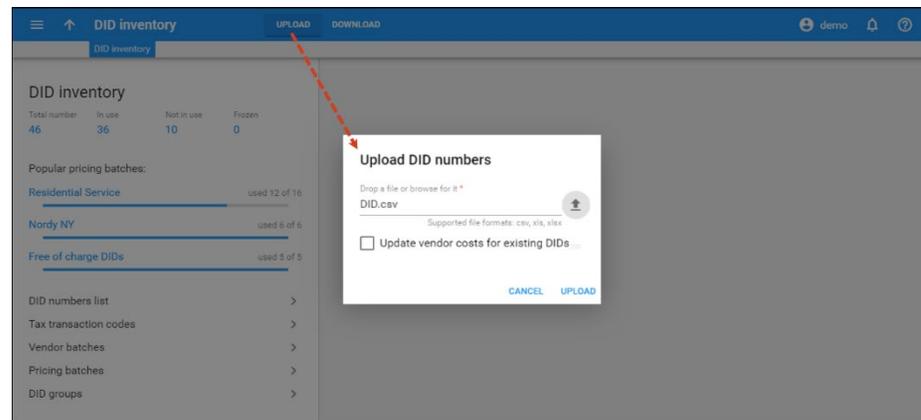
When you as the service provider obtain DID numbers from a DID provider, you are usually charged an activation fee and then monthly recurring costs for all of the DID numbers. Therefore, when provisioning DID numbers to your customers, you apply an additional markup to these charges in order to compensate for your costs and earn some profit. The parameters of the markup are defined in pricing batches.

Please note that activation costs do not increase the DID provider's balance (a DID provider is provisioned as a vendor in PortaBilling®). Monthly recurring costs are recalculated according to the vendor's billing period and applied for all obtained DID numbers regardless of whether or not they are assigned to customers. These costs are what increase the vendor's balance.

DID costs management

When your DID provider changes their pricing and sends you an updated DID list, all you have to do is upload it to the DID inventory.

To update the vendor costs in the DID inventory, select the **Modify data of existing DIDs** option during the DID upload. The new DID numbers from the .csv file, if any, will simply be uploaded to their respective vendor batches in the DID inventory.



Upon upload, the new vendor costs override previous costs for all DID numbers regardless of their status (Free, In Use, etc.). You will receive a notification stating the number of both new and updated DIDs that have been uploaded to the DID inventory.

DID fees (calculated as the sum of vendor costs and DID markup) are also modified automatically. Thus, customers using DID numbers are charged updated recurring fees. Their activation fees will not be recalculated.

Billing for DID number usage

Every DID number belongs to a pricing batch. The pricing parameters of this batch, if any, define which charges will be applied to a customer when a DID number is assigned to them.

This allows you to set up different pricing for different DID number provisioning (e.g. local or international ones).

Consider the following example:

A customer is assigned a US number with \$5 in activation costs and \$3 in recurring costs from the “Domestic numbers” pricing batch. The markup for this batch is defined as follows: a \$1 additional activation fee and a \$2 additional recurring fee. The customer is charged \$6 (\$5+\$1) upon DID activation and a \$5 (\$3+\$2) monthly recurring fee recalculated according to the customer’s billing period (i.e. if a customer has a semimonthly billing period, he is charged \$2.5 per billing period).

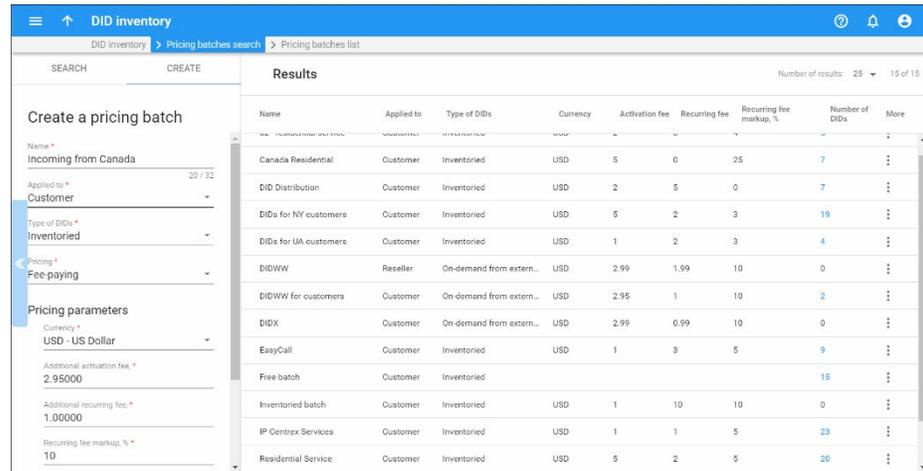
Then the customer is assigned a Chinese DID number with \$15 in activation costs and \$5 in recurring costs from the “International

numbers” pricing batch. This batch has the following DID markup defined: a \$5 additional activation fee and \$4 additional monthly recurring fee. The customer is now charged \$20 (\$15+\$5) for DID activation and \$9 (\$5+\$4) for the monthly recurring fee.

DID pricing parameters

DID pricing is defined by the following parameters:

- **Additional activation fee** is a configurable fixed amount that is charged when a new number is provisioned, in addition to the DID provider’s activation fee.
- **Additional recurring fee** is a configurable fixed amount that is always included in the recurring (monthly) charge.
- **Recurring fee markup** is a configurable increase of the DID cost (charged by the DID provider) included in the recurring charge in addition to the actual DID cost. For instance, when defining the ability to dynamically assign US numbers from DIDWW, we can specify that we would like to include a 10% markup on the DID provider rates. Thus, if a customer assigns phone number 12065551234 (DIDWW rate \$5/month), the end user will be charged \$5.50/month.



The screenshot shows the 'DID inventory' interface. On the left, there is a 'Create a pricing batch' form with the following fields: Name (Incoming from Canada), Applied to (Customer), Type of DIDs (Inventoried), Pricing (Fee-paying), Currency (USD - US Dollar), Additional activation fee (2.95000), Additional recurring fee (1.00000), and Recurring fee markup (10). On the right, there is a 'Results' table with 15 rows of pricing data.

Name	Applied to	Type of DIDs	Currency	Activation fee	Recurring fee	Recurring fee markup, %	Number of DIDs	More
Canada Residential	Customer	Inventoried	USD	5	0	25	7	⋮
DID Distribution	Customer	Inventoried	USD	2	5	0	7	⋮
DIDs for NY customers	Customer	Inventoried	USD	5	2	3	19	⋮
DIDs for IA customers	Customer	Inventoried	USD	1	2	3	4	⋮
DIDWW	Reseller	On-demand from extern...	USD	2.99	1.99	10	0	⋮
DIDWW for customers	Customer	On-demand from extern...	USD	2.95	1	10	2	⋮
DIDX	Customer	On-demand from extern...	USD	2.99	0.99	10	0	⋮
EasyCall	Customer	Inventoried	USD	1	3	5	9	⋮
Free batch	Customer	Inventoried					15	⋮
Inventoried batch	Customer	Inventoried	USD	1	10	10	0	⋮
IP Centrex Services	Customer	Inventoried	USD	1	1	5	23	⋮
Residential Service	Customer	Inventoried	USD	5	2	5	20	⋮

As a result, the total charge for the DID would be:

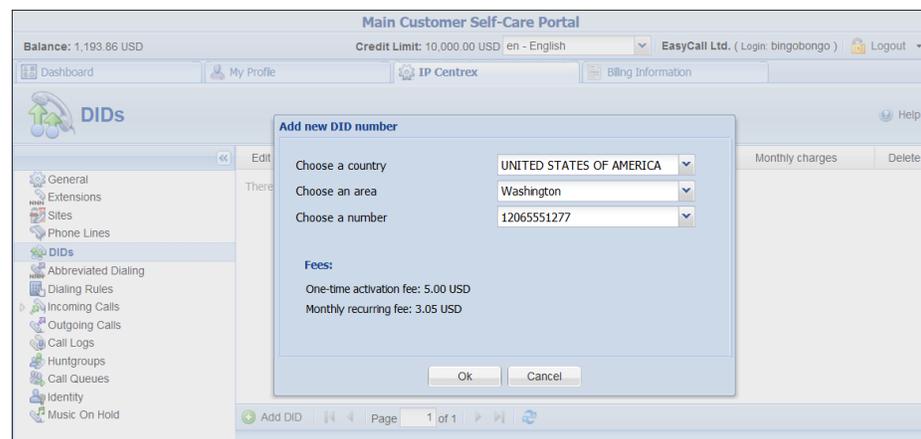
- **At activation** – DID provider activation charge + Additional activation fee.
- **Monthly** – DID provider monthly charge + the markup, consisting of the Additional recurring fee and the Recurring fee markup. This formula is likely to result in some “inconvenient” numbers, such as 3.289, which are usually rounded up.

Rounding

In order to obtain “marketable” amounts to be presented on the customer’s invoice, specify the rounding pattern for the charged amounts. This will pre-set the last digits in the amount.

For instance, in a situation where number 12065551234 (which has a \$1.00 activation fee and a \$3.00 monthly fee charged by the DID provider) is provisioned by an end user, the DID pricing setup would be as follows:

- \$2.95 for additional activation fee,
- \$1 for additional recurring fee,
- 10% for recurring fee markup,
- and 2 decimals pattern for rounding (this parameter means that a charge amount will be rounded up to two decimal places).



The screenshot displays the 'Main Customer Self-Care Portal' interface. At the top, it shows the user's balance (1,193.86 USD), credit limit (10,000.00 USD), and language (en - English). The user is logged in as 'EasyCall Ltd. (Login: bingobongo)'. The main content area is titled 'DIDs' and features a sidebar menu with options like General, Extensions, Sites, Phone Lines, DIDs, Abbreviated Dialing, Dialing Rules, Incoming Calls, Outgoing Calls, Call Logs, Huntgroups, Call Queues, Identity, and Music On Hold. A modal dialog box titled 'Add new DID number' is open, allowing the user to select a country (UNITED STATES OF AMERICA), an area (Washington), and a number (12065551277). The dialog also displays the associated fees: a one-time activation fee of 5.00 USD and a monthly recurring fee of 3.05 USD. The dialog has 'Ok' and 'Cancel' buttons.

If a customer has a shorter billing period, e.g. weekly, a proportion of the DID provider’s monthly fees will be used to calculate the customer’s charge.

NOTE: Please keep in mind that if you are going to *provisionally* terminate a customer, this customer won’t be charged for any DID numbers assigned to them. At the same time, the DID provider will still charge you a fee (e.g. \$5/month) because these DID numbers remain allocated to your network until the customer is *permanently* terminated.

Completely free of charge DID numbers

When you provide virtual DID numbers for internal use (e.g. for SIP calls among accounts within your network), you still want to see which of them are in use and which are available. The best way to do this is by using the DID inventory. However, since these DIDs are virtual, they are, as a rule, provisioned free of charge and incur no charges to customers.

Therefore, to manage such internal DID numbers differently the concept of *completely free of charge* DID has been introduced. These are the DID

numbers that have no costs or charges associated with them and are therefore free of charge for every entity they are provisioned to.

When an administrator assigns a completely free DID number to a pricing batch and then provisions the DID number to a customer, no charges apply.

Since there are no recurring costs associated with the DID number, the vendor is not billed either. Thus, both customer and vendor balances remain unchanged since no xDRs are created in the database.

Consider the following example:

You operate with thousands of end users. The services you provide allow end users to make and receive calls within your network. Since end users are not available from outside of your network, they do not require geographic DID numbers. However, to identify them within your network, you provision them virtual DID numbers. For this, you use the pool of virtual DID numbers 87899800000-87899989999.

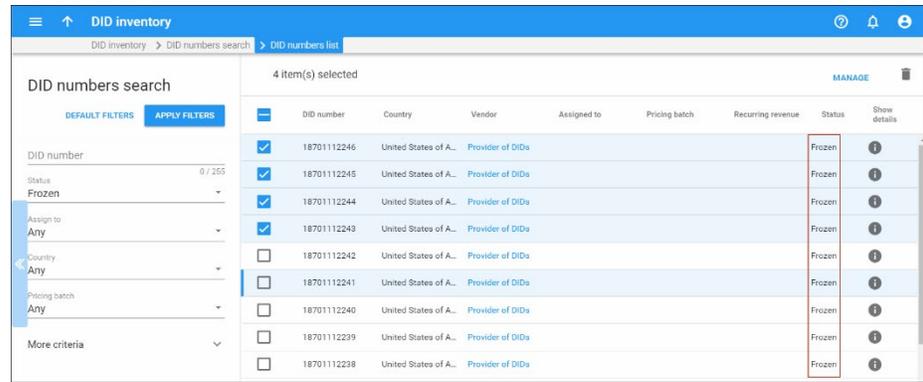
Since these are your own internal numbers, you bear no costs and apply no charges to customers for their usage. Therefore you have them as completely free of charge in the DID inventory.

When a customer signs up for a product directly or via a reseller, no costs for DID provisioning apply to him or her.

Frozen DIDs and DID pool

It may be that you do not want to immediately provision DID numbers acquired from a DID provider (e.g. you plan to release a new product next month and therefore need to keep the DIDs until then). In this case you need to prevent the DID numbers from being used by your customers and resellers.

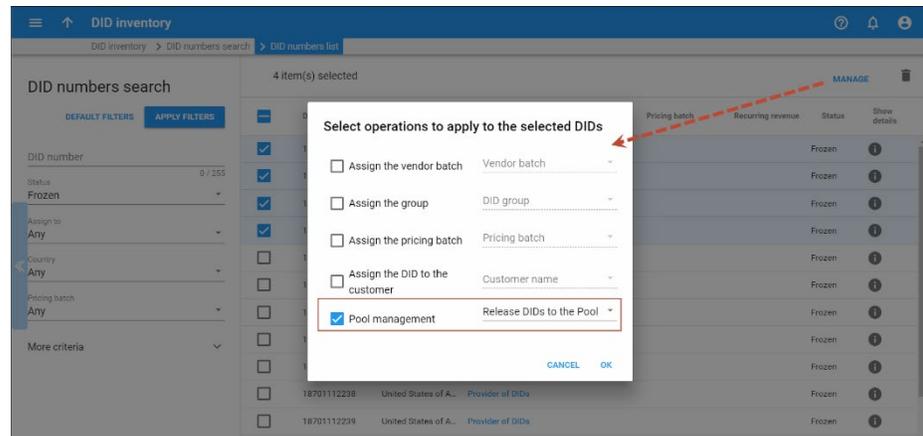
In order to keep the DIDs unavailable for usage, they are uploaded to the DID inventory under the **Frozen** status. When frozen, DIDs are displayed on the administrator's DID inventory page only.



The screenshot shows the 'DID inventory' interface with a search filter set to 'Frozen'. A table lists 8 DID numbers, with the first four selected. The 'Status' column for these four items is highlighted in red.

Selected	DID number	Country	Vendor	Assigned to	Pricing batch	Recurring revenue	Status	Show details
<input checked="" type="checkbox"/>	18701112246	United States of A...	Provider of DID's				Frozen	
<input checked="" type="checkbox"/>	18701112245	United States of A...	Provider of DID's				Frozen	
<input checked="" type="checkbox"/>	18701112244	United States of A...	Provider of DID's				Frozen	
<input checked="" type="checkbox"/>	18701112243	United States of A...	Provider of DID's				Frozen	
<input type="checkbox"/>	18701112242	United States of A...	Provider of DID's				Frozen	
<input type="checkbox"/>	18701112241	United States of A...	Provider of DID's				Frozen	
<input type="checkbox"/>	18701112240	United States of A...	Provider of DID's				Frozen	
<input type="checkbox"/>	18701112239	United States of A...	Provider of DID's				Frozen	

To allow resellers to provision DID numbers, the administrator simply releases them to the pool of available DID's.



The screenshot shows a modal dialog box titled 'Select operations to apply to the selected DID's'. The 'Pool management' option is selected, with 'Release DID's to the Pool' chosen from the dropdown menu. A red dashed arrow points from the dialog to the 'Status' column of the table in the background screenshot.

The DID pool is the tool that allows you to control the DID numbers that will be provisioned by your resellers. It contains the numbers that are:

- not frozen,
- not assigned to pricing batches, and
- not used as an account ID, account alias, or an access number.



In order to retrieve DID numbers from the pool, resellers must have pricing batches with a DID markup assigned to them, since this markup is what they are charged for DID number usage.

Consider the following example:

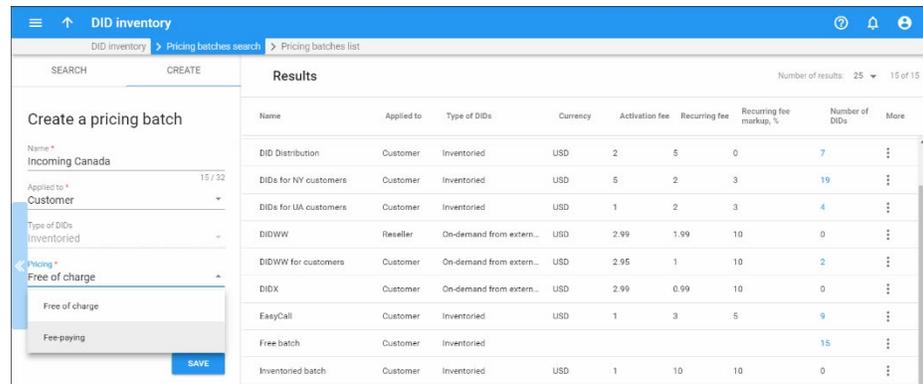
You are releasing a new product in two weeks – the IP Centrex solution for your corporate customers. For this you have acquired the following range of DID numbers: 12045556500 – 12045556599. You must be sure that these DID numbers remain reserved for provisioning until the product is released. Therefore you keep them frozen in the DID inventory.

When the product is released, you unfreeze the DID numbers by releasing them to the DID pool and then distribute them to respective pricing batches.

Or, if you want your resellers to provision DIDs along with the new product, leave some of these DIDs in the DID pool.

Assigning DID numbers to pricing batches

You can provision DID numbers to your customers either with your markup or free of charge. The type of pricing batch a DID number is assigned to is what determines whether it is free or not. The pricing batch type is selected during the batch creation.

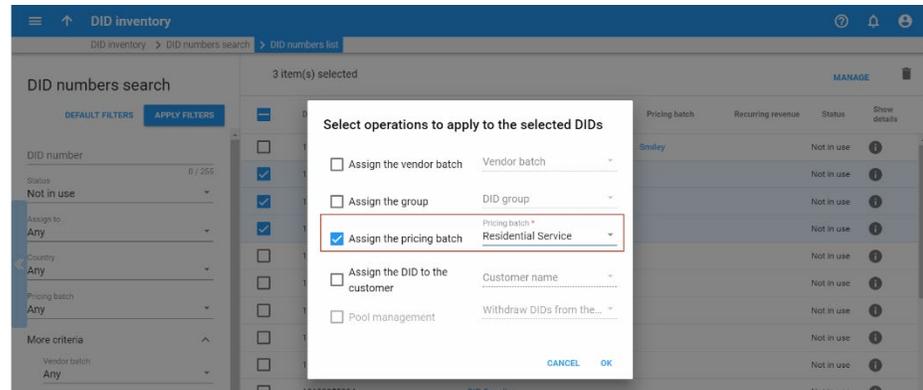


The screenshot shows the 'DID inventory' interface. On the left, there is a 'Create a pricing batch' form with the following fields: Name (Incoming Canada), Applied to (Customer), Type of DIDs (Inventoried), and Pricing (Free of charge). A 'SAVE' button is at the bottom. On the right, there is a 'Results' table with columns: Name, Applied to, Type of DIDs, Currency, Activation fee, Recurring fee, Recurring fee markup, %, Number of DIDs, and More. The table contains several rows of pricing batch data.

Name	Applied to	Type of DIDs	Currency	Activation fee	Recurring fee	Recurring fee markup, %	Number of DIDs	More
DID Distribution	Customer	Inventoried	USD	2	5	0	7	⋮
DIDs for NY customers	Customer	Inventoried	USD	5	2	3	19	⋮
DIDs for IA customers	Customer	Inventoried	USD	1	2	3	4	⋮
DIDWW	Reseller	On-demand from extern...	USD	2.99	1.99	10	0	⋮
DIDWW for customers	Customer	On-demand from extern...	USD	2.95	1	10	2	⋮
DIDX	Customer	On-demand from extern...	USD	2.99	0.99	10	0	⋮
EasyCall	Customer	Inventoried	USD	1	3	5	9	⋮
Free batch	Customer	Inventoried					15	⋮
Inventoried batch	Customer	Inventoried	USD	1	10	10	0	⋮



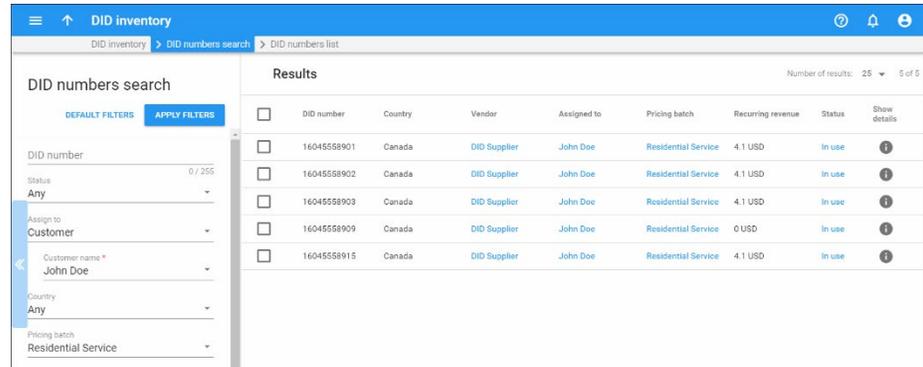
To properly manage DID numbers and control your costs and revenue, it is important to assign each DID number to an appropriate pricing batch prior to allocating this number to a customer.



The screenshot shows the 'DID numbers list' interface. A modal window titled 'Select operations to apply to the selected DIDs' is open. The modal has several options with checkboxes: 'Assign the vendor batch' (Vendor batch), 'Assign the group' (DID group), 'Assign the pricing batch' (Residential Service), 'Assign the DID to the customer' (Customer name), and 'Pool management' (Withdraw DIDs from the...). The 'Assign the pricing batch' option is checked and highlighted with a red box. The background shows a list of DID numbers with columns for Pricing batch, Recurring revenue, Status, and Show details.

When a DID number is assigned to a customer from a pricing batch with a DID markup, the markup is applied to the customer.

Pricing batches reflect a DID markup specific for either an administrator or a reseller. An administrator does not see any pricing batches created by resellers because those contain resellers' markups.



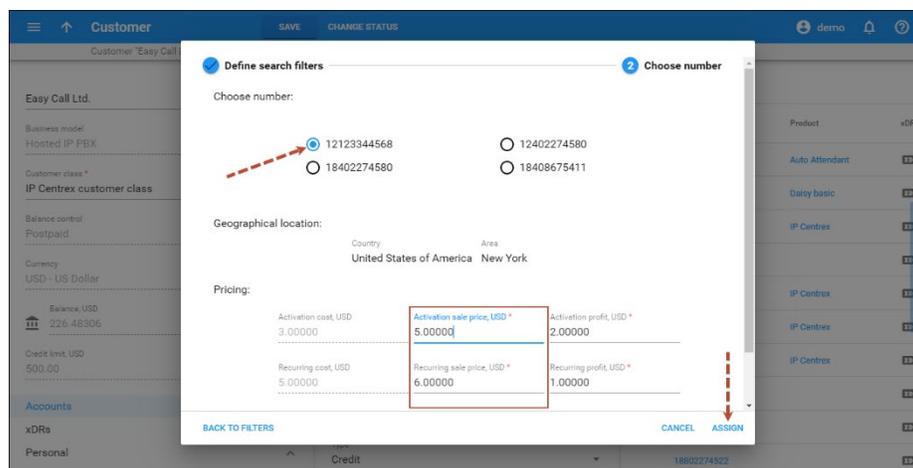
<input type="checkbox"/>	DID number	Country	Vendor	Assigned to	Pricing batch	Recurring revenue	Status	Show details
<input type="checkbox"/>	16045558901	Canada	DID Supplier	John Doe	Residential Service	4.1 USD	In use	
<input type="checkbox"/>	16045558902	Canada	DID Supplier	John Doe	Residential Service	4.1 USD	In use	
<input type="checkbox"/>	16045558903	Canada	DID Supplier	John Doe	Residential Service	4.1 USD	In use	
<input type="checkbox"/>	16045558909	Canada	DID Supplier	John Doe	Residential Service	0 USD	In use	
<input type="checkbox"/>	16045558915	Canada	DID Supplier	John Doe	Residential Service	4.1 USD	In use	

When added to a free pricing batch, a DID number becomes free of charge. This means that no charges apply to either a customer or a reseller when a DID number from a free pricing batch is assigned to them (e.g. DID number 16045552879 has a \$10 activation cost associated with it and is assigned to the free pricing batch. Customer John Doe signs up for the EasyCall product and is provided with the number 16045552879. He is charged nothing for the DID number activation yet the administrator now owes \$10 to the DID provider).

NOTE: You can move DIDs from the pricing batches they are currently assigned to to other pricing batches. If a DID number is not in use, it can be reassigned to any pricing batch, regardless of whether it is free or has a DID markup. However, if a DID number is allocated to a customer or a reseller, it can only be reassigned to the same type of pricing batch (i.e. to free pricing batches or to pricing batches with DID markup).

Individual pricing for DID numbers

Your customers may ask for a personal price when purchasing a DID number. Instead of creating a special pricing batch with this new markup and moving the DID numbers to it, now modify the existing one per DID. When you create an account, click the **DID** button, choose a particular DID number and specify the individual pricing for it.



A DID number with modified fees is automatically marked as having **Individual pricing**.

DID groups

To easily manage a large number of DID numbers, administrators may use **DID groups** – a tool that allows arranging DID numbers into categories (for example, by geographic location).

Consider the following example. The service provider wants to distribute a group of North American DID numbers free of charge, a group of Canadian and Western European DID numbers at \$3 per month and a group of DID numbers from other countries at a 20% markup.

For this scenario, an administrator creates three DID groups and distributes the DID numbers among them. The administrator then creates three pricing batches with different pricing schemes and assigns the corresponding pricing batch to each DID group. The DID numbers are now ready to be distributed to customers.

Note that separating DID numbers into groups is an optional step which is solely intended to make DID number management more flexible and convenient.

DID provisioning to customers

Depending upon your business model, you can provision DID numbers to your customers so that:

- a customer is charged the same price for using different DID numbers (\$5 monthly for US and UK DID numbers). This is defined by the pricing parameters of the pricing batch directly assigned to the customer;

- a customer is charged different prices for using different DID numbers (e.g. \$5 monthly for the US DID and \$8 monthly for the UK DID number). This is defined by the pricing parameters of the pricing batch that a particular DID belongs to.

Additionally, you can define whether customers will be charged as soon as the DID numbers are assigned to them (and not yet provisioned as customer accounts) or only after the DID numbers have been provisioned to the customers as accounts or account aliases.

NOTE: When provisioning a DID number from a pricing batch to a customer and if this customer has another pricing batch specifically assigned to him, the charges defined in the customer's pricing batch will be ignored. Instead, the charges defined in the pricing batch that a DID number belongs to will apply.

If customers provision DID numbers from external vendors (i.e. by using on-demand DID provisioning), it is required that the pricing batch having the DID markup be assigned to them because they are billed according to the pricing parameters defined in this specific pricing batch.

A DID number explicitly assigned to a customer and used as an account or account alias is released from this customer once the alias or the account is terminated or changed. Although it remains within the respective pricing batch, it has free status within the DID inventory and is available for further provisioning.

Customers can keep the DIDs assigned to them (e.g. to create other aliases). To make this happen, a customer clicks Edit and unassigns the DID from the alias. This helps you fulfill your contractual obligations if customers have purchased DIDs for the long term.

However, if the customer is terminated, the DID number is released from the customer and has free status within the DID inventory.

Please refer to the *DID Provisioning* section of the **Unified PortaSwitch Handbook Collection** for detailed instructions of ways to provision DID numbers to customers.

DID number self-provisioning

Customers can self-provision the service provider's DID numbers. For customers' convenience, DID numbers are grouped by country and country area.

Consider the following example:

The service provider purchased a set of DID numbers and wants them to be made available to their customers with minimal administrative assistance.

An administrator defines the DID markups in pricing batches and distributes the DIDs to them.

If a customer needs some DID numbers for their business, they log in to the customer self-care interface and allocate some of the available DID numbers. The customer is then charged the price specified for the selected DIDs.

If the customer decides that they no longer need the DID number, they simply delete it from their list of allocated DIDs, the number is released from this customer and the customer is no longer charged for the use of it. The DID number can now be provisioned by other customers.

DID	Country	Assigned to	Monthly charges	Delete
12065551236	UNITED STATES OF AMERICA Washington	Unassigned	1.10 USD	X

Note that when the customer deletes a provisioned DID number, the pricing batch it previously belonged to is preserved. So the DID number returns to the DID inventory and is once again available for all customers (regardless of whether or not they are allowed to provision DID numbers themselves).

If prior to deletion, the DID number was assigned to an account as an alias, this alias is removed in PortaBilling®.

This allows customers to provision DID numbers by themselves, without administrative assistance, thus decreasing an administrator's workload.

Reseller DID provisioning

Access to the DID pool allows resellers to use as many DID numbers as they need and then provision these numbers to their customers. This expands the service provider's DID provisioning network and boosts sales.

To access the DID pool, resellers must have pricing batches assigned to them.

NOTE: Subresellers access the DID pool only if they and their top-level resellers have pricing batches assigned to them.

Then, they can perform the following actions when provisioning DID numbers via the DID inventory.

These actions include:

- Definition of the pricing strategy for DID number provisioning;
- Distribution of DIDs retrieved from the DID pool to pricing batches with respective DID markups;
- Assigning DID numbers to subcustomers;
- Releasing DID numbers to the DID pool.



DID Inventory											
DID Number		Customer	Group	Pricing Batch	Country	Status	Search		Advanced Search		
<input type="checkbox"/>	0445652155					Free	11.99 USD	6.01 USD			Calls to Ukraine
<input type="checkbox"/>	0445652159					Free	11.99 USD	6.01 USD			Calls to Ukraine
<input type="checkbox"/>	12124254782		UNITED STATES OF AMERICA			Free for 54 days	7.00 USD	8.05 USD			United States of America
<input type="checkbox"/>	12124254790		UNITED STATES OF AMERICA			Free for 54 days	7.00 USD	8.05 USD			United States of America
<input type="checkbox"/>	1604123610		CANADA	ABC Shuttle Ltd.		Frozen	12.00 USD	7.04 USD			Canada
<input type="checkbox"/>	1604123675		CANADA			Free for 54 days	12.00 USD	7.04 USD			Canada
<input type="checkbox"/>	1604123676		CANADA	ABC Shuttle Ltd.		Frozen	12.00 USD	7.04 USD			Canada
<input type="checkbox"/>	1604123677		CANADA			Free for 54 days	12.00 USD	7.04 USD			Canada

The DID pricing strategy definition

As a preparatory step, a reseller defines the revenue he wants to receive as the DID markup after provisioning this or that DID number within a pricing batch. To associate a DID markup with a DID number, add the number to the corresponding pricing batch.

Another pricing strategy is to define the same DID markup for any type of DID number provisioned to the same customer (e.g. 10% to be added to the DID recurring cost). This is done by assigning a default pricing batch to a customer. Then, after a DID number is provisioned to the customer, add it to their pricing batch and the defined DID markup will be applied. Such a model is frequently used when DID costs are at first unknown (on-demand DID provisioning).

Retrieval of DID numbers from the DID pool

To manage DID numbers, a reseller retrieves them from the DID pool by assigning them to respective pricing batches. Once a DID number is retrieved from the DID pool, it is owned by the reseller and is therefore unavailable for provisioning by other resellers. Though the DIDs may not

yet be provisioned to customers, charges for them are already applied to their owners: the resellers.

NOTE: If a top-level reseller exceeds their credit limit, their subreseller will not be able to retrieve a DID number from the DID pool.

DID provisioning to customers

When a DID number is provisioned to a subcustomer, the charges are calculated for each party involved and are distributed as follows:

- A top-level reseller is charged a DID number cost plus additional fees defined in the pricing batch assigned to this reseller;
- A subreseller is charged the top-level reseller's costs plus additional fees as defined in the pricing batch assigned to this subreseller;
- A subcustomer is charged the subreseller's costs plus the DID markup as defined by the subreseller in the pricing batch associated with this number.

Release of DID numbers to the DID pool

Any DID number managed by a reseller can be released by this reseller back to the DID pool if the DID number is not used as an account ID, account alias, or an access number. To do this a reseller must select the **Release DIDs to the pool** option and click **Apply**. While in the DID pool, the DID number is again visible to all other resellers and is available for provisioning.

Please note that if a reseller un-assigns a DID number from a pricing batch in the DID inventory, this number is not released to the DID pool. The reseller still owns this number and still bears the charges for its usage.

A reseller can provision a DID number to a customer either from a pricing batch or via a pricing batch that's specifically assigned to this customer. Thus, when a reseller deletes the alias or changes the account / alias ID, the assigned DID number is released as follows:



- In case a DID number is provisioned to a customer from a pricing batch, it is released back to this pricing batch. This means that the reseller continues to bear the charges for this DID, though the customer is no longer charged for it;
- In case a DID number is provisioned to a customer via a pricing batch specifically assigned to this customer, a DID number is automatically released to the pool. This means neither the reseller nor the customer owns this DID number and therefore no charges apply for the DID usage.

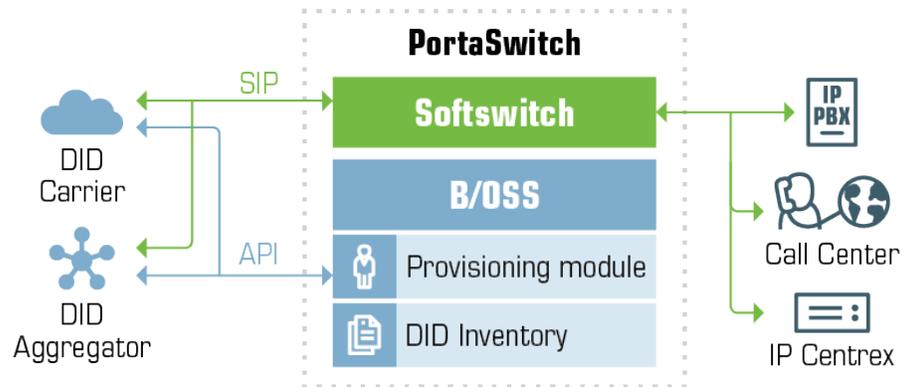
The reseller does not receive any notification about a DID release when they manage accounts or aliases.

The flexibility of DID provisioning for resellers reduces the load on your administrators and boosts sales.

On-demand DID provisioning and billing

While the traditional DID provisioning method is still used when an ITSP obtains a large pool of phone numbers and then one by one, assigns numbers from it to individual customers, the appearance of “whole world a-la carte” providers such as Voxbone, DIDWW or DIDX has changed the preferred way of delivering extra DID numbers from foreign countries as a value-added service. It is impractical to purchase blocks of numbers from every single country in the world in the hope that some customer is going to buy them. DID providers usually charge a monthly fee for each allocated number, so this mode will quickly become hard to sustain financially.

On-demand DID provisioning enables the end user to browse the complete inventory of DID or toll-free numbers available on the DID vendor side, online. When he picks a number, the sign-up process on the PortaBilling® side will reserve and allocate that number on the DID provider side and provision it inside of PortaBilling® (e.g. as an alias to an existing account). When the customer decides that he no longer needs the number (or changes the number to something else), the reverse procedure will happen – the number will be released on the DID provider side and the account/alias will be removed in PortaBilling®. This gives you the ability to offer a wide variety of different DID numbers to the end user and only be charged by the DID provider for numbers actually used by customers.



Interconnection with DID providers

All major DID vendors offer some type of API that can be used to query, reserve and release DID numbers. As APIs are different for different DID providers, a system of plug-in modules has been implemented (similar to those used for payment systems) that will easily allow the addition of new providers in the future.

The initial version of the DID provisioning module includes integration with DIDWW (<http://www.didww.com>) and DIDX (<https://www.didx.net>) providers – and more providers will be added soon based on requests from our customers.

Combined DID provisioning

To generate additional revenue you can grant customers with on-demand DID provisioning access both to your DID number pool and that of other external DID providers (DIDWW, DIDX, etc.).

Consider the following example:

EasyCall has a pool of DID numbers for Canada. The service provider SmartNet (with whom EasyCall exchanges DID numbers) owns DID numbers for the New York area.

EasyCall enables on-demand DID provisioning for their customer John Doe. So now John is able to self-provision Canadian and New York DIDs via his customer self-care interface.

To enable this functionality, the administrator adds the `[DID]MixedSources=1` option for the **WebCustom** group on the Configuration server.

DID exchange among service providers

To grow the DID provisioning network, you can now grant access to your DID inventory to other ITSPs – DID buyers – via the API, and allow them to provision DID numbers from you on-demand. This enables you to become an active player at the wholesale DID market.

This is how it works:

On-demand DID provisioning for DID buyers is done via the DID provisioning interface, using the API. You configure the DIDPool server – client communication on the Configuration server, and then define the access to your DID inventory using the credentials of your DID provisioning customer. This customer will accumulate charges for all DID numbers provisioned from your system, therefore they must have the DID pricing batch assigned.

Your DID buyers consequently define you as their external DID vendor and thereby configure on-demand DID provisioning in their systems. Thus, your PortaSwitch® operates as the DIDPool server while their systems function as DIDPool clients.

Of course, the routing among your systems must also be configured.

As before, the DID management itself is done within the DID inventory. When your DID buyer's customer provisions a number from the DID pool, the following occurs:

- The DID number in the DIDPool server system is assigned to your DID provisioning customer and acquires In Use status within the DID inventory.
- The DID number appears in the DIDPool client system as being assigned to that customer.

The charges for DID provisioning and usage are calculated in both the DIDPool server and DIDPool client systems as follows:

- The DID provisioning customer in the DIDPool server system is charged according to the DID pricing batch assigned to them.
- The customer in the DIDPool client system is charged for DID number usage according to the markup defined in their DID pricing batch.
- If a DID number is no longer required, it is released from the customer and returned back to your DID pool.

Consider the following example:

The service provider, EasyCall, owns a pool of 10 million DID numbers. SmartNet is another service provider who exchanges voice traffic with EasyCall. SmartNet wants to introduce residential VoIP services to their system and thus requires DID numbers.

EasyCall allows SmartNet to provision DIDs from them via their customer DID Distribution. To do this, EasyCall provides SmartNet with DID Distribution's credentials.

SmartNet sets up DID provisioning in their system and provisions DID Distribution as the external DID provider. Now they can allocate DID numbers from EasyCall like they can from any other external provider.

When SmartNet no longer needs DID numbers and releases them to the DID pool, the numbers are returned to the EasyCall DID pool with the Free (not in use) status.

Thus, the ability to distribute DID numbers to other service providers and charge for their usage extends your clientele and provides vast opportunities for wholesale DID trade.

Registration for PINless dialing service via SMS

This is an extension to the “classic” PINless dialing service, in which the user is identified by his caller ID (ANI/CLI), maintains a balance in PortaSwitch and is able to call any destination with charges applied to his balance. Traditionally, the registration for PINless dialing service is performed by the end user via the purchase of a top-up card and then calling the special IVR. When the user’s balance is low, he needs to buy more top-up cards or pay his bill online. The registration via SMS (simply send an SMS message to a certain number) is an easier alternative. A simplified registration process and further management of the account make this service especially convenient for users of mobile networks. The service provider has two options for collecting funds from the end user:

- Send a request to the mobile operator that the end user is connected to for a balance transfer. This option must be supported by the mobile carrier – its advantage is that any “ordinary” number can be used by the customer for sending an SMS.
- Establish a premium number where the end user can send an SMS. Just like calls to premium numbers are charged differently than normal calls, messages to premium numbers are charged according to the higher rate (set by the service provider who owns the number). So for each message sent by the end user, the mobile carrier will collect the money from the end user and then pass it to you as the owner of the premium number.

To start using the service, the end user simply sends an SMS with the message SUBSCRIBE to the registration number. Upon receiving the message, the SMS processing module in PortaSwitch automatically provisions his account with the CLI/ANI of the mobile phone as the authentication ID; requests a balance transfer from the mobile operator and then adds these funds to the customer’s balance in PortaSwitch so they can be used again in the future. End user will receive an SMS that confirms the service activation and provides important service information, such as the number to call for PINless dialing or a username/password for accessing the web self-care. Once that occurs, the end user can start using the service by dialing the access number and then entering the final destination number. The entire registration process is very easy to perform and can be completed in just a few seconds.

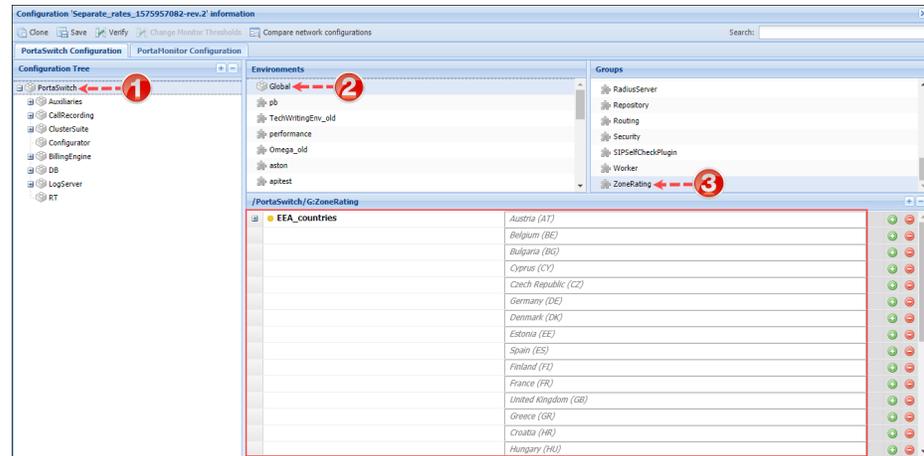
PortaSwitch can also perform an automated balance top-up: when the balance of the created account goes below a specified threshold, the system initiates another balance transfer from the user’s balance on the mobile network. The telecom client can disable this automatic recharge by sending an SMS to the SMS access number with the message STOP.

Origin-based pricing for European destinations

Some European carriers apply separate charges for calls from within / outside the European Economic Area (EEA). Call costs are determined both by call origin (the country the call is made from) and destination (the country where it terminates.)

You can charge wholesale customers differently for calls from within / outside the EEA, too. You define separate rates for the same destination for calls made from both inside and outside the EEA. Then for each call to a European destination, PortaBilling® checks whether the CLI belongs to the EEA countries and applies the corresponding rate.

The list of EEA countries is predefined in PortaBilling®. An administrator can adjust the default list of EEA countries in the **ZoneRating.EEA_countries** option on the Configuration server.



This is how it works:

Let’s say your wholesale customer, EasyCall, sends calls from all over the world to Malta. You route these calls to your vendor GlobalNet for termination.

GlobalNet charges you separately for calls to Malta made from within / outside the EEA as follows:

Destination	Rate within EEA EUR/min	Rate outside EEA EUR/min
356	0.1	0.5

You charge EasyCall as follows:

Destination	Rate within EEA EUR/min	Rate outside EEA EUR/min
356	0.3	0.8

Let's consider these examples:

1. GlobalNet terminates a call from Italy to Malta.

Since Italy belongs to the EEA countries, the rate within the EEA is applied to this call. Thus, GlobalNet charges you 0.1 EUR per minute while you charge EasyCall 0.3 EUR per minute for this call.

2. GlobalNet terminates a call from the USA to Malta.

The USA destination does not belong to the EEA list so GlobalNet applies the rate for outside the EEA for this call – 0.5 EUR per minute, and you charge EasyCall 0.8 EUR per minute.

Private calls processing

By default, the CLI is provided in the 'From' SIP header. If an end user requests that their phone number be hidden from the party called (the **Hide CLI** feature is enabled), the 'From' header will contain "Anonymous" instead of the CLI. The actual CLI is stored in the PAI (P-Asserted Identity) SIP header and that's the header PortaBilling® uses to match the rate.

Your vendors can also deliver private incoming calls wherein the CLI is stored in the PAI SIP header. For PortaBilling® to match the rates given by the PAI, these conditions must be met:

- The call must arrive from a trusted network (the **Caller identity** service feature is set to **Supply** for the incoming vendor connection).
- The account is configured to accept caller's information: the **CLI trust** and **Supply caller identity** options are enabled.

Calls with overridden identity

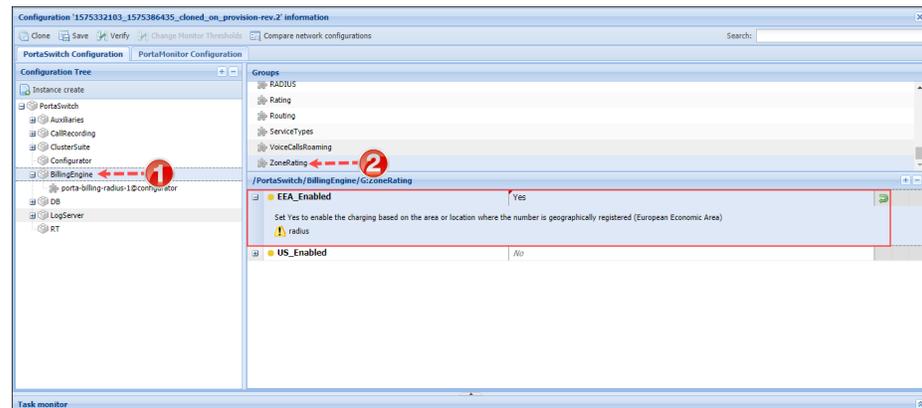
A user can request that their identity be overridden for outgoing calls. For example, a company can insist that their employees use the company's main phone number 340755511340 for all calls. When such a call arrives to PortaSIP®, it includes the PAI header with the actual CLI (e.g.12065552211). PortaBilling® uses the PAI header value to match the rate.

NOTE: Rate match by the PAI SIP header is supported only if calls arrive via PortaSIP® for processing. If you use another termination gateway, PortaBilling® attempts to retrieve the CLI from the 'From' SIP header to match the rate. And if the CLI cannot be retrieved, PortaBilling® matches the rate by CLD.

Configuration:

To enable separate rates for calls from within / outside the EEA follow these steps:

1. Set the **EEA_Enabled** option to **Yes** for the **ZoneRating** group on the Configuration server.



2. Add rates for calls within / outside the EEA to your customer tariff. We recommend uploading these rates from either a .csv or .xls file for wholesale customers. Please refer to the **Rate Import** handbook for complete details.

The screenshot shows the 'Add rate' configuration page for a 'Wholesale tariff'. The left sidebar contains a menu with 'Rates' highlighted. The main form area has the following fields and options:

- Destination:** 356
- Rating mode:**
 - Flat rate
 - Separate peak/off-peak rate
 - Separate rates within/outside EEA
- First interval, seconds:** 1
- Next interval, seconds:** 1
- Rate outside EEA:**
 - First price, EUR/minute: 0.80000
 - Next price, EUR/minute: 0.80000
- Rate within EEA:**
 - First price, EUR/minute: 0.40000
 - Next price, EUR/minute: 0.40000

Note that setting peak / off-peak prices is not supported for the **Separate rates within / outside EEA** rating mode.

OLI-based billing

Vendors that provide toll-free numbers to service providers can charge a different price for calls that originate from a payphone or a paystation.

With PortaSwitch®, service providers can take call origination into account when charging their customers and calculating calls' cost on the incoming connections (Calls from Vendor via SIP and Calls from Vendor via PSTN).

Call origination type is determined based on the Origination Line Information (OLI) parameter. The OLI is transmitted over PSTN networks as signals. There is no common standard for carrying OLI over SIP, but one popular way is to use the `isup-oli` SIP URI parameter and include it in the SIP INVITE message, notably in the `From` header.

For example, for calls made from a payphone, the `From` header may look like the following:

```
From: "18005551155"
<sip:18005551155@111.11.11.11:5060;isup-
oli=27>;tag=gK037dd60b
```

PortaSIP® extracts this value and assigns it to the `Originating-Line-Info` RADIUS attribute which will be then passed to PortaBilling®:

```
NAS-IP-Address = '111.11.11.11'
User-Name = '18005551155'
...
Originating-Line-Info = '27'
...
```

Using this value, the billing engine seeks a rating entry that can match it, and if one is found, it both charges for the call and calculates the call cost accordingly.

OLI (Info Digits)

OLI digit pairs and their descriptions are represented in the following table:

Value	Description
0	pots
1	multiparty line
2	ANI failure
6	Station Level Rating
7	special operator handling required
8	inter-LATA restricted
10	test call
20	AIOD-listed DN sent
23	coin or non coin on calls using database access
24	800 service call
25	800 service call from a pay station
27	payphone using coin control signaling

29	prison or inmate service
30	intercept (blank)
31	intercept (trouble)
32	intercept (regular)
34	telco operator handled call
36	CPE
52	OUTWATS
60	TRS call from unrestricted line
61	wireless or cellular PCS (type 1)
62	wireless or cellular PCS (type 2)
63	wireless or cellular PCS (roaming)
66	TRS call frm restricted line
67	TRS call from restricted line
68	inter-LATA restricted hotel
70	private paystations
78	inter-LATA restricted coinless
93	private virtual network

6. MNO / MVNO billing

Integration with mobile networks

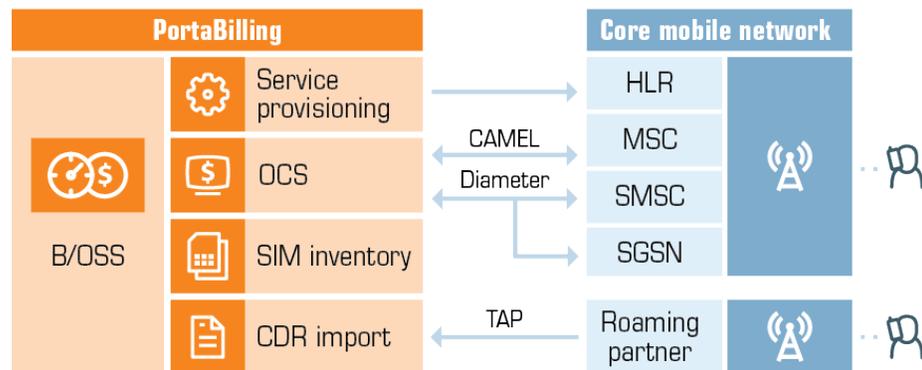
For a service provider planning to become a mobile operator, there are two alternative options for entering the mobile market. The first one, and perhaps the most straightforward one, is through acquiring a spectrum license and deploying your own mobile network.

The other market entrance option is to become a virtual mobile network operator (MVNO) and utilize the network infrastructure of an existing MNO (mobile network operator). This model allows virtual operators to offer to their customers a full range of mobile services without having to make significant investments to build and maintain a mobile network.

PortaBilling® can serve as the B / OSS and OCS for any of these solutions and integrate with 2.5 / 3G and 4G mobile networks.

Integration with 2.5 / 3G networks

The integration process between PortaBilling® and 2.5 / 3G networks includes the following major aspects:



Customer data provisioning

When an MNO signs up a mobile end user (subscriber) for a service, the subscriber is provisioned as a customer with one or more accounts in PortaBilling®. Each account represents an actual phone line with an assigned phone number (MSISDN) and is associated with a SIM card (identified by IMSI).

This subscriber must be recognized in the mobile operator’s network in order for them to be granted access to the service. To do this, PortaBilling® communicates with the HLR (Home Location Register) mobile network component.

The HLR is a centralized repository that stores information about all of an MNO's subscribers. Once an account is added and activated in PortaBilling®, the system connects to the HLR and creates an entry with the account's IMSI (SIM card) and MSISDN (phone number).

Real-time charging

Real-time authorization and rating is an important requirement for providing prepaid mobile services. PortaBilling® fully supports this, allowing real-time charging for services such as voice calls, SMS messages or Internet use.

In 2.5 and 3G networks, the CAMEL protocol is the de-facto standard for real-time charging. Therefore, PortaBilling® interoperates with mobile operator networks via a CAMEL gateway (GW) that converts the CAMEL protocol to the Diameter protocol and vice versa.

The mobile network component that PortaBilling® interoperates with for real-time authorization and rating is the MSC (Mobile Switching Center).

When a subscriber makes a call, MSC performs call authorization in PortaBilling® to verify that the subscriber has a sufficient balance and is permitted to call that destination. Since MSC uses the CAMEL protocol, its call authorization request is first converted to Diameter Ro and then delivered to PortaBilling®. PortaBilling® locks funds to cover the initial period of the call and sends confirmation to MSC. When that period is over, MSC sends a request to lock funds for the subsequent period, etc. Once the call is over, MSC produces a CDR record and sends it to PortaBilling®.

In the message exchange between MSC and PortaBilling®, all of the protocol conversions – from CAMEL to Diameter and from Diameter to CAMEL – are done using the CAMEL GW.

A similar procedure is followed when a customer sends an SMS. In this case, PortaBilling® interoperates with the SMSC (Short Message Service Center) – a mobile network component that is responsible for delivering SMS messages.

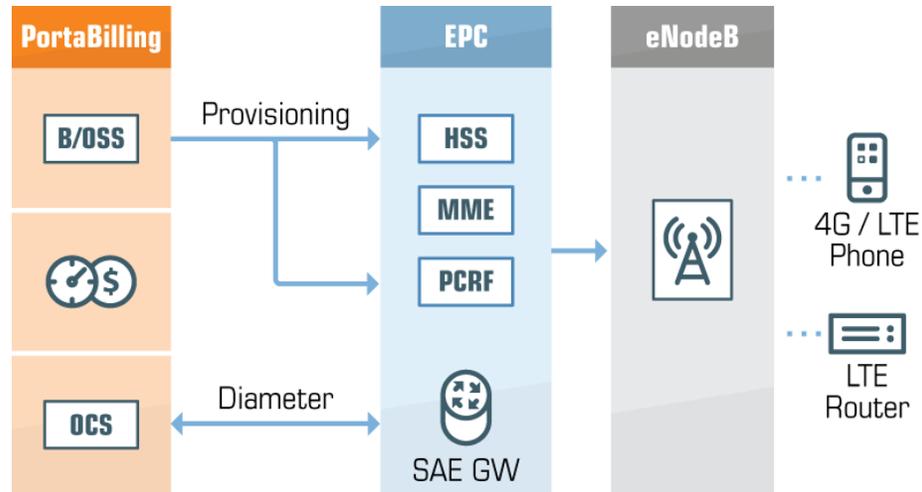
Roaming

PortaBilling® permits MNOs to include roaming with their service offering.

When an MNO subscriber roams to a place where their MNO is not covered, they are automatically connected to another MNO – a roaming partner – without any interruption in service.

The roaming partner generates CDRs with service usage information in the roaming area and sends them to the subscriber's MNO in TAP files. TAP is a commonly used standard for roaming data exchange among mobile operators. The MNO then parses these TAP files to distinguish the roaming CDRs that belong to their subscriber and sends them to PortaBilling®. PortaBilling® processes those CDRs and charges the subscriber for all of the roaming services provided.

Integration with 4G / LTE networks



A similar level of integration is done with 4G / LTE networks.

To provision customer's data, PortaBilling® communicates with the following components:

- **HSS (Home Subscriber Server)** – This is a database that contains information about all of the MNO's subscribers. Once a subscriber is configured in PortaBilling® and assigned a SIM card and a desired product, the subscriber's information is nearly instantly provisioned in the HSS. Now the subscriber is ready to register on the mobile network.
- **PCRF (Policy and Charging Rules Function)** – This component provides policy decisions for mobile network subscribers. Such policies are usually related to QoS parameters and charging rules. PortaBilling® provisions customer data to PCRF to allocate the correct policy / QoS parameters to the subscriber.

For real-time session authorization and rating, PortaBilling® communicates with SAE-GW (System Architecture Evolution Gateway) via Diameter (Gy) for Internet service provisioning and via Diameter (Ro) for voice calls and SMS delivery.

Once the subscriber's session starts, SAE-GW performs real-time authorization in PortaBilling® and locks the funds to cover the initial period of the session. When the session is over, SAE-GW transmits the information about the transferred data volume to PortaBilling®. PortaBilling® then charges the subscriber for the session.

Mobile Virtual Network Operator

To penetrate the mobile services market, there is no need to purchase expensive hardware, build your own tower network, or (in some cases) even get a license from the regulatory authorities. As a Mobile Virtual Network Operator (MVNO), you can utilize an existing mobile network to sell your products and bill your customers directly. PortaBilling serves as the B/OSS and OCS for such an MVNO.

Service providers can choose between two MVNO models: one with or one without real-time authorization and rating.

In the MVNO model *with* real-time authorization and rating, PortaBilling® interacts with the mobile operator's MSC or SAE-GW (depending on the type of hosting mobile network) as described above.

In the MVNO model *without* real-time authorization and rating, the mobile operator that hosts the MVNO will need to receive information about accounts and their status and deliver information about completed calls to the MVNO. Thus, there are two major aspects of the interaction between the MVNO's billing and provisioning core (PortaBilling) and that of the hosting mobile operator's network:

Account management

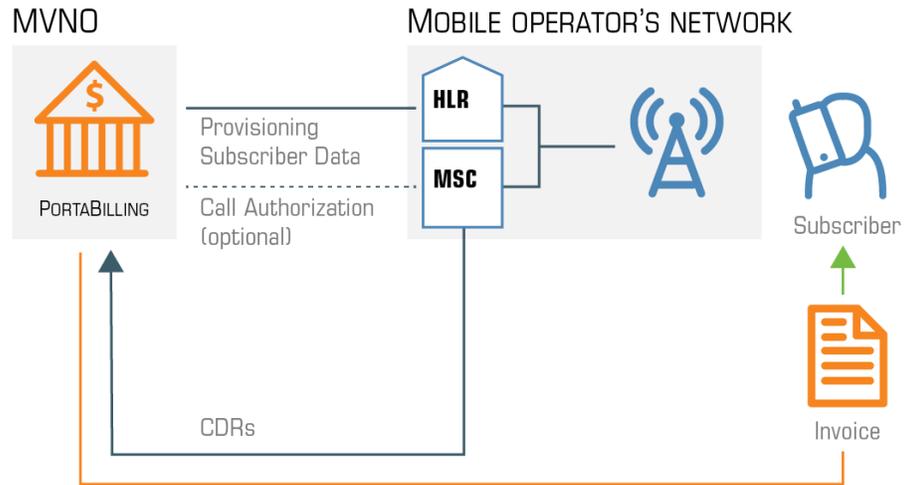
When an account/customer is created by the MVNO, it should be provisioned to the HLR (Home Location Register) of the host mobile operator's system. Account blocking / unblocking, deletion, and so on should also be mapped to the mobile operator.

This is achieved by using External System Provisioning Framework (ESPF) event mapping. When a change event occurs in PortaBilling (e.g. creation of a new account, product change, or suspension of a customer due to non-payment), it is queued and an integration plug-in is invoked. The application code inside the plug-in then delivers the information to the mobile operator's network using the required protocol (e.g. text files or XML API calls) and communication method.

Call data retrieval and processing

To process call data for an MVNO's customers, the mobile operator periodically delivers files to the MVNO containing the CDRs for all calls

made by the MVNO's customers during a certain time interval. On the PortaBilling side, a regular **xDR import** is then performed.



Billing customers and sending out invoices is done directly by PortaBilling®, without involving the host mobile operator.

SMS routing with real-time HLR lookup

In mobile networks, message routing is based on a combination of the mobile country code (MCC) and the mobile network code (MNC). The MCC / MNC pair identifies the mobile carrier within the network and is used to build routing lists. Mobile carriers also operate with MCC / MNC codes to generate SMS pricing.

To facilitate effective SMS delivery, message routing with real-time HLR lookup is introduced in PortaSwitch®. This means that PortaSwitch® requests subscriber information from the HLR (Home Location Registry) in real-time. It obtains the MCC / MNC codes of the destination number and then uses them for SMS delivery and billing.

Thus, two aspects of SMS provisioning are considered in the scope of this document: MCC / MNC rate management and HLR lookup.

MCC / MNC rate management

The majority of mobile carriers operate with MCC / MNC codes for destinations and rates. Therefore, for convenience in working with these carriers, define destinations in the E.212-MCC-MNC format in PortaBilling®.

SEARCH	CREATE	Results				
Destination search		<input type="checkbox"/>	MCC	MNC	Country	Network
DEFAULT FILTERS APPLY FILTERS Format: E212 MCC: Any MNC: Any Country: Any Network:		<input type="checkbox"/>	202	02	Greece	Cosmote
		<input type="checkbox"/>	402	17	Bhutan	Bhutan Telecom Ltd (BTL)
		<input type="checkbox"/>	738	01	Guyana	DigiCel
		<input type="checkbox"/>	232	15	Austria	T-Mobile/Telering
		<input type="checkbox"/>	750	001	Falkland Islands	Touch
		<input type="checkbox"/>	289	88	Georgia	A-Mobile
		<input type="checkbox"/>	289	77	Georgia	
		<input type="checkbox"/>	289	65	Georgia	

You can then upload your rates for these destinations in the tariffs for Messaging services and test the routing by using the SMS routing simulation.

NOTE: To test message routing, make sure that HLR lookup is enabled on the Configuration server.

HLR lookup

HLR lookup serves two goals: to identify which mobile network the destination number belongs to and to perform correct SMS routing and billing.

To identify which mobile network the destination number belongs to and to which mobile carrier route the message, PortaSIP® performs HLR lookup. This is done by sending an HLR lookup request to the external provider – www.hlr-lookups.com. This is the default HLR lookup provider, however, other providers can be integrated upon request.

Upon successful HLR lookup, PortaSIP® receives the MCC and MNC codes for the destination number. The MCC / MNC pair defines which mobile carrier handles this number, even if it was ported from another carrier.

The MCC / MNC codes are passed to PortaBilling® to compute a routing list and calculate vendor costs and user charges.

Configuration

To provide SMS services using real-time HLR lookup, set up an account within www.hlr-lookups.com and configure the **HLRNumberLookup** module on the Configuration server.

Please refer to the **Unified PortaSwitch Handbook** collection for a detailed configuration description.

NOTE: Delays in message delivery to a carrier are possible.



If you enable HLR lookup, only vendors that support the E.212-MCC-MNC format will participate in the SMS routing. For customer / reseller tariffs you can define rates in either E.164 or E.212 format. However, to properly calculate charges, reseller and customer tariffs must have the same format.

The support of MCC / MNC codes simplifies rate management and the configuration of SMS providers within PortaBilling®. Real-time HLR lookup assures the accuracy of SMS routing and billing. Consequently, you boost your competitiveness within the mobile telecommunications market and attract more customers.

Real-time billing for roaming services

To remain competitive in the market, wireless operators must ensure that their subscribers can enjoy roaming services. Roaming service charges are higher than domestic ones; therefore, real-time billing is crucial for saving subscribers from overspending.

PortaBilling® supports real-time billing for roaming data, voice calls and SMS services.

Billing for roaming data

To bill for roaming data usage, PortaBilling® selects a tariff rate based on the MCC-MNC codes that identify the subscriber's visited network. Please read more in the Billing for the LTE service in roaming section.

Billing for roaming voice calls and SMS

To bill for roaming voice calls and SMS, PortaBilling® uses the identification of both the visited network and the roaming zone to match the tariff in the product and produce the charges.

All networks are grouped into roaming zones. Each roaming zone has its own rates for outgoing and incoming calls and SMSs defined in tariffs.

Every network has an application server that processes calls. There is also the SMSC to process SMSs. The E.164 numbers of these servers are used to identify the network, map it to the roaming zone and then select the corresponding roaming tariff.

The charging algorithm is this:

When a user makes / receives a call from their home network, PortaBilling® uses the “home” tariffs with domestic rates to charge them.

For roaming calls, PortaBilling®:

1. identifies the roaming zone based on the application server's E.164 number;
2. selects the corresponding roaming tariff, and
3. charges the user for outgoing calls based on the destination number. Incoming calls are charged the same price, determined by roaming zone.

SMSs are processed in the same way. The difference is that PortaBilling® uses the serving SMSC's E.164 number to identify the roaming zone.

Mapping of roaming zones

PortaBilling® uses the “Roaming Zones” destination group set to identify roaming zones. Every destination group within the set corresponds to a roaming zone and includes corresponding prefixes. The home network is excluded from the destination group set.

Workflow for voice calls

Your product rating list must include these usage charges records:

- all roaming tariffs with the OUTGOING.<destination_group_name> and INCOMING.<destination_group_name> access codes.
- “home” tariffs with the default OUTGOING and INCOMING access codes.

When a user makes a call, PortaBilling® receives the Diameter Ro request from the mobile core. This request contains the Application-Provided-Called-Party-Address AVP (attribute-value pair) with the application server's E.164 number. PortaBilling® uses this number to look up the longest prefix match in the “Roaming Zones” destination group set.

The destination group name where the match is found serves to formulate the access code as either OUTGOING.<destination_group_name> or INCOMING.<destination_group_name> and match the tariff in the product rating list.

If no match is found, then this is a call within the home network. PortaBilling® charges such calls using the default OUTGOING and INCOMING access codes.

Usage scenario

Let's say you are the MVNO in the UK. Your host MNO has roaming agreements with T-Mobile Polska in Poland and Sprint in the USA. T-Mobile Polska belongs to roaming zone 1, Sprint – to roaming zone 2.

Your administrator creates the Roaming Zones destination group set. It includes:

- A destination group Zone1 with a 48602 prefix in it;
- A destination group Zone2 with a 1913 prefix in it.

Your pricing is:

From	Home country, GBP	Zone1, GBP	Zone2, GBP
To			
Home country, GBP	0.10	0.30	0.70
Zone1, GBP	0.60	0.80	1.25
Zone2, GBP	1.0	1.20	0.90
Incoming calls while roaming, GBP	0	0.45	0.75

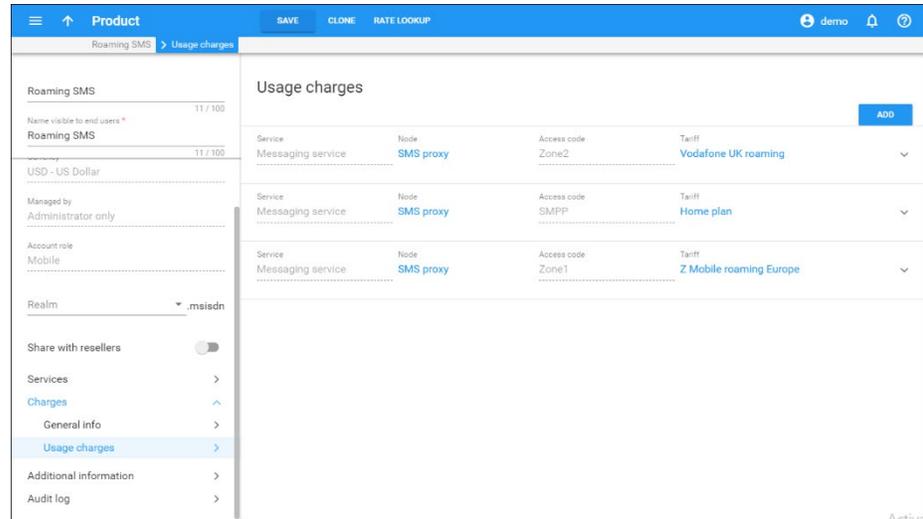
John Doe is your subscriber. He travels to Poland and makes a 5-minute call to the USA while there.

PortaBilling® receives the Diameter Ro request with 486025552237 stored in the Application-Provided-Called-Party-Address AVP. PortaBilling® looks up this number among the prefixes in the Roaming Zones destination group set and finds the match in destination group “Zone1.” PortaBilling® uses this name to match the usage charges record with the OUTGOING.Zone1 access code in John’s product and John is therefore charged GBP6.20 (GBP1.20 * 5).

While still in Poland, John receives an incoming call. PortaBilling® performs the same destination lookup and finds the Zone1 destination group. PortaBilling® matches the usage charges record using the INCOMING.Zone1 access code and charges John GBP 0.45 / min.

Workflow for SMS

PortaBilling® receives the Diameter Ro request where the serving SMSC’s E.164 number is stored in the Originator-SCCP-Address AVP. PortaBilling® uses this number to look up the longest prefix match among the destination groups in the “Roaming Zones” destination group set. The destination group serves to identify the user’s roaming zone. This destination group name matches the access code defined for the usage charges record in the product configuration. Thus, PortaBilling® uses this name to select the roaming tariff and charge the subscriber.



If no match is found, then this is an SMS sent from the home network. PortaBilling® charges for this SMS using the default SMPP access code.

Usage scenario

Let’s say that John Doe from Canada travels to the UK. He is connected to Vodafone UK. Vodafone UK belongs to roaming zone 2 where SMSs to Canada cost \$0.03. The “Roaming Zones” destination group set in PortaBilling® contains the destination group “Zone2” which includes 44385 prefix.

John sends an SMS to his wife in Canada. PortaBilling® receives the Diameter Ro request with the number 443855551555 in the Originator-SCCP-Address AVP. PortaBilling® finds the longest match for 443855551555 in the destination group “Zone2.” PortaBilling® searches for the usage charges record with the Zone2 access code in John’s product configuration and selects the corresponding roaming tariff. As a result, John is charged \$0.03 for an SMS.

Benefits

With this functionality, the benefits for mobile subscribers include:

- the ability to use services regardless of their location;
- a clear view of one’s own expenses and therefore, better control.

Wireless operators, in their turn, benefit from being able to:

- deliver mobile services to subscribers in both their home network and when they visit other countries;
- reduce the risk of revenue loss by charging subscribers in real time; and
- configure roaming charges in a single place and in a user-friendly fashion.

Supported solutions

Real-time billing for calls based on the roaming country's rates is tested with the Protei CAMEL gateway and available out of the box. Billing for roaming SMSs is supported with Protei SMSC.

Submit your request to support@portaone.com to perform interoperability tests with your equipment.

Implementation specifics for roaming calls and SMS billing:

- Real-time billing for roaming calls and SMSs is enabled via a special module for Diameter. To enable it, enter the `roamingzone=Porta::Extra::RoamingAccessCode` string in the **Custom.LoadModules** option on the Configuration server.
- You must create the destination group set with the name “Roaming Zones” in PortaBilling®. This is to identify the user's roaming zone and consequently, select a roaming tariff.
- Access codes in product rating lists must correspond to the destination group names in the “Roaming Zones” destination group set.

Known limitations:

- PortaBilling® charges only for outgoing SMSs.
- The user is charged for sending SMSs, regardless of whether the SMS is delivered or not.

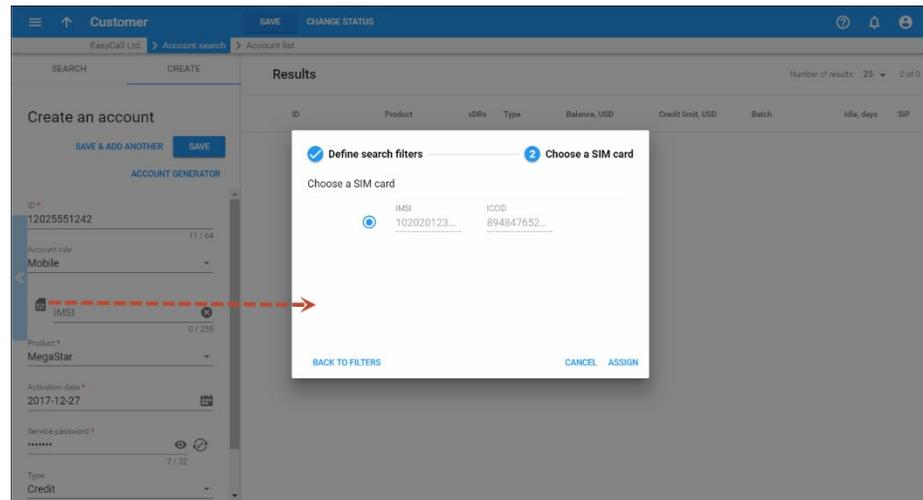
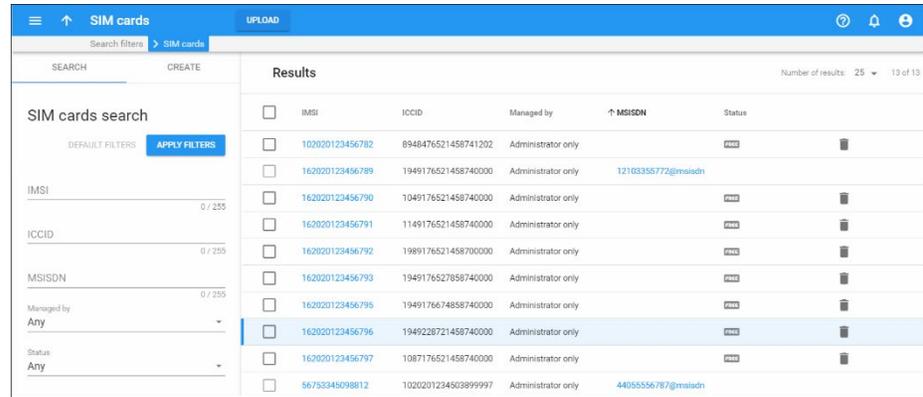
SIM card inventory

The SIM card inventory is a flexible tool that enables the MVNOs and LTE service providers to bulk upload SIM cards via the administrator web interface and therefore, easily manage them.

The SIM card inventory is used for adding, viewing, removing and assigning SIM cards to accounts. Its convenient search function makes it possible to quickly find the desired record by simply inputting a key word.

SIM cards can be added to a SIM card inventory either manually, i.e. one at a time, or can be bulk uploaded from a .csv file. Find out more about SIM card uploads and the .csv file structure in the [PortaBilling® Help](#).

Once a SIM card is added to a SIM card inventory, its status is *Free*. This means that this SIM card is not yet assigned to any account. When a SIM card is assigned, its status is *In Use*. Note that a SIM card having an *In Use* status cannot be removed from the SIM card inventory until it is first unassigned.



When a SIM card is assigned to an account, it acquires the MSISDN value that equals the *account ID* of the linked account. Therefore, both the account and the assigned SIM card are now connected and PortaBilling® can automatically provision the information into HSS / HLR by using the External System Provisioning Framework.

With the increasing need to manage and maintain large numbers of SIM cards and provision them into external systems, this SIM card inventory is becoming a great tool for LTE service providers and MVNOs.

Mobile account batch generation for SIM card distribution

You can either generate mobile accounts manually or in batches via the Account generator. Batch generation saves time, as you can create a large number of accounts at once. Each account in a batch is assigned a mobile number (MSISDN) and a SIM card (IMSI) from the corresponding inventory.

Once a batch of accounts is generated, the accounts' data is automatically provisioned to the host mobile network (i.e., HSS and/or PCRF), if PortaBilling® is integrated with the mobile carrier's core.

The screenshot shows the 'Create an account' interface with an 'Account generation' modal open. The modal is titled 'Account generation' and has a sidebar with 'Initial configuration', 'Optional configuration', 'Service', and 'Summary'. The 'Initial configuration' section is expanded and shows the following fields:

- Account role * (Mobile)
- Type * (Debit)
- Product * (Super One)
- Quantity * (100)
- Batch (Add a new batch)
- New batch name * (10%SIMcards)

At the bottom right of the modal are 'CANCEL' and 'GENERATE' buttons. The background form has an error message: 'Error: Please enter a value' under the 'ID' field.

Batch generation reduces administrators' manual work and also enables mobile operators to involve distributors in selling ready-to-use SIM cards to a large number of end users.

For example, a new distributor John Smith wants to sell your SIM cards in a local supermarket chain. You need to provide John with 100 SIM cards. To configure this batch, your administrator opens the Account generator to create 100 inactive debit accounts with the Mobile role. The administrator sets the initial account balance to \$10 and chooses the starting MSISDN and IMSI number for this batch. Note that this data should be previously uploaded to the DID/MSISDN and SIM card inventories. To assign accounts to John, the administrator chooses his name in the Distributor drop-down list.

NOTE: Only debit accounts can be assigned to a distributor.

Account generation

Initial configuration @msisdn, 100, 10%SIMcards > Assign MSISDN sequentially starting with
 The first one available

Optional configuration >

Service > Assign IMSI sequentially starting with
 The first one available

Summary >

Service password *
 Auto generated

Customer site
 None

Distributor
 John Smith

Life cycle >

E-commerce enabled

CANCEL GENERATE

Now, the new accounts provided to the distributor are ready to use. When John Smith sells a SIM card, he activates the account via the web self-care interface in one click. The end user can start using the service right away.

The distributor’s commission for activation and further account top up per every SIM card is automatically reflected in PortaBilling®. Refer to the **Distributor** section for more details.

Your xDR History

xDRs Revert

Customer	John Smith	Total Transactions	3
From Date	2020-05-17 17:58:21	Credits	0.00000 USD
To Date	2020-05-18 17:58:25	Charges	18.00000 USD
Service (s)	All	Recurring Charges	0.00000 USD
		Recurring Credits	0.00000 USD
		Usage Charges	0.00000 USD
		Usage Credits	0.00000 USD
		Payments	0.00000 USD
		Refunds	0.00000 USD
		Taxes	0.00000 USD
		Total	18.00000 USD

Credits / Adjustments				
Description	Comment	Date/Time	Amount, USD	Revert
Distributor: Transfer	C.Mobile services provisioning	2020-05-17 20:57:22	0.00000	<input type="checkbox"/>
Distributor: New product sale	Charge for account:12060898684@msisdn	2020-05-17 20:23:25	9.00000	<input type="checkbox"/>
Distributor: New product sale	Charge for account:12060898683@msisdn	2020-05-17 20:22:50	9.00000	<input type="checkbox"/>
Subtotal			18.00000	

Thus, mobile account batch generation enables mobile operators to expand their sales network by selling SIM cards through distributors.

Local number portability and other special cases

While call processing (routing, billing) based on a phone number match to a phone prefix (when each phone prefix defines a group of phone numbers such as 420 – Czech Republic; 4202 – Czech Republic, Prague; 420604 – Czech Republic, T-Mobile and so on) is used in the majority of telecommunications services, in some cases it is not applicable.

Local number portability

This is a very common example of a situation where the rate or call routing cannot be determined simply by looking at the phone number itself. Local number portability (LNP) is widespread throughout Europe, the US and other countries around the world. It means that when a user migrates from one telco (fixed or mobile operator) to another, he is able to take his existing phone number with him. So although his phone number looks as if it belongs to one company (e.g. mobile operator A), in reality his calls should be routed to another company (mobile operator B). Most importantly, anyone making a phone call to this number should be charged according to the termination rates defined by B.

The only way to determine whether a number has been ported or not is to check it against the ported numbers database. These databases differ for each individual country, and have a variety of formats and access API structures. However, PortaBilling® offers a flexible method for handling LNP, using a system of number translation plug-ins that satisfy the requirements valid in various countries.

A basic plug-in for LNP (where all ported numbers are imported into a local database on the PortaBilling® server) is supplied along with the system. When a call processing request arrives to PortaBilling®, the billing engine checks if the original number dialed by the user is found on the list of ported numbers. If there is no match, call setup proceeds as usual. However, let us assume that a customer has dialed a number that originally belonged to the mobile operator Telenor, but was recently ported to another operator, Netcom. In this case, there will be an entry in the LNP table, and the ID of the phone number's new owner will be extracted (4792, in this example). This ID could be a symbolic name (e.g. Netcom), but usually it is more convenient to use a phone prefix or phone number belonging to that operator. The following will then happen:

- The billing engine looks up routing to a destination identical to the number owner's ID (4792, in this case), and the call is routed directly to the Netcom network.

- The same ID is used to calculate your termination costs and apply charges to your customer, i.e. he will be charged according to Netcom's rates.
- The originally dialed phone number is stored in the CDR record, but is linked to the destination used for charging the call. So the end user will see "4791555123 – Norway, Netcom mobile" in the CDR report.

If a different way of extracting information from the LNP database is required for a given country (for instance, real-time lookup in a central database), you can create your own LNP plug-in and use it with PortaBilling®.

The Number_Portability table in PortaBilling® is described in the [APPENDIX D. The Number_Portability Table Description](#) section.

LRN (Local Routing Number) dipping

In the United States and Canada, the only way to determine whether a number has been ported is to subscribe to an LRN dipping service that looks up ported numbers in real-time from a remote database.

Since PortaSwitch® is connected with Telnix, an LRN dipping service provider, it provides this opportunity to determine whether a number has been ported – and it routes the call to the most suitable and most economic vendor.

When a call processing request arrives at PortaBilling®, the following happens:

1. PortaBilling® checks the destination number and if necessary sends a request to the LRN service provider (e.g. PortaBilling® will only check US & Canadian numbers and skip the rest).
2. If the number is not ported, the provider's response indicates that the call should be routed, based on the original dialed number. However, let us assume that a customer dials a number (for example, 12063131234), that originally belonged to the mobile operator T-Mobile, but was recently ported to AT&T. In this case, the provider returns the LRN number (e.g. 12065549080) and PortaBilling® uses *it* to route the call directly to the AT&T network.
3. When the call finishes, the same LRN number is used to calculate the termination costs and then the charges are applied to the customer – in this case, according to AT&T's rates. The number originally dialed is stored in the CDR record, but is linked to the destination used for charging the call. So the end user will see "12063131234 – USA, AT&T" in the CDR report.

- In addition, to improve speed and flow, PortaBilling® will *temporarily* add an entry for this number and its corresponding LRN number to the list of ported numbers.

Since PortaSwitch® is already interconnected with Telnyx, it can offer a flexible plug-in solution that allows for easy integration with other LRN service providers.

Porting numbers from / to PortaBilling®

A service provider can now perform the following operations directly from the unified PortaBilling® web interface:

- Port a number to PortaBilling® from another telecom (port-in);
- Port a number from PortaBilling® to another telecom (port-out).

This feature is convenient for customers who want to change telephone providers but keep their original numbers. This framework provides the option to port numbers via Neustar Inc. – a major player in the number portability market.

In order to port-in a number, a PortaBilling® user sends a request via the **Port in** button on the **Edit Account** page. In the request dialog window the following parameters must be mentioned:

- The phone number that the customer wants to use;
- The service password or PIN number of a customer's account within the previous service provider's system;
- The date in the PortaBilling® user's time zone by which this request must be completed;
- The customer's account within the previous service provider's system.

The screenshot shows a web interface for editing an account. A modal dialog box titled "Port in number" is open, containing the following fields and values:

- Number to port: 1604043734
- PIN for the ported number: 2541575
- Due date: 2018-11-17 05:23
- OSP internal account: 1000999555

At the bottom of the dialog box, there are two buttons: "CANCEL" and "PORT IN".

PortaBilling® sends the porting request to Neustar and waits for confirmation. Neustar requests the number from the telecom and when confirmation is received, the number is replaced in PortaBilling® and the customer can receive calls at the new number.

When PortaBilling® receives a port-out request, it searches the account with the requested number. If the account is found, PortaBilling® terminates it and informs Neustar about the number’s availability. Neustar then sends this information to the telecom that requested the number. So now that the ported number belongs to this telecom, calls are routed to the new telecom’s network.

NOTE: PortaBilling® automatically provides the requested number (port-out); administrators cannot affect the process.

Let’s consider the following example:

John Doe is a customer of X-Telecom and has the number 1212555123. Due to questionable connection quality, he decides to switch from X-Telecom to GlobalNet Telecom. He calls GlobalNet Telecom, asks to use their services, but declares his wish to keep his old number. The GlobalNet Telecom operator informs John Doe that the service will be available starting January 1st. Then the operator creates an 1818123555 account for John Doe in GlobalNet Telecom, makes the porting request via PortaBilling® to Neustar, and when the request is confirmed, the system replaces 1818123555 with 1212555123. So while John Doe continues to use his old 1212555123 number, the calls are now charged according to GlobalNet Telecom’s tariffs.

The PortaBilling® administrator may monitor all requests from his system (the **Port-in** panel) and to his system (**Port-out** panel) by clicking the **Infrastructure**, then **Inventory** and **Number porting** buttons on the **navigation** menu.

PORT-IN		PORT-OUT				
		Number of results: 25 / 6 of 6				
<input type="checkbox"/>	Status	Account ID	Porting number	Due date	Request details	Adjust
<input type="checkbox"/>	New	16041235064	124-097-8536	2018-11-18 16:00:00		
<input type="checkbox"/>	New	1778551200	160-404-3734	2018-11-18 16:00:00		
<input type="checkbox"/>	New	1604556920	443-423-6236	2018-11-18 16:00:00		
<input type="checkbox"/>	New	1870112235	128-973-7578	2018-11-18 16:00:00		
<input type="checkbox"/>	New	17096545678	160-927-8367	2018-11-18 16:00:00		
<input type="checkbox"/>	New	15198833154	207-468-3487	2018-11-18 16:00:00		



NOTE: The **Port in** button is only available when **NumberPorting** is enabled on the Configuration server. To use Neustar for number portability, **NumberPorting_Neustar** must also be enabled.

PortaBilling® can receive number porting requests via the HTTP / HTTPS protocols.

Thus, this new feature makes it possible for customers to use their original numbers, regardless of the telecom. Number portability works

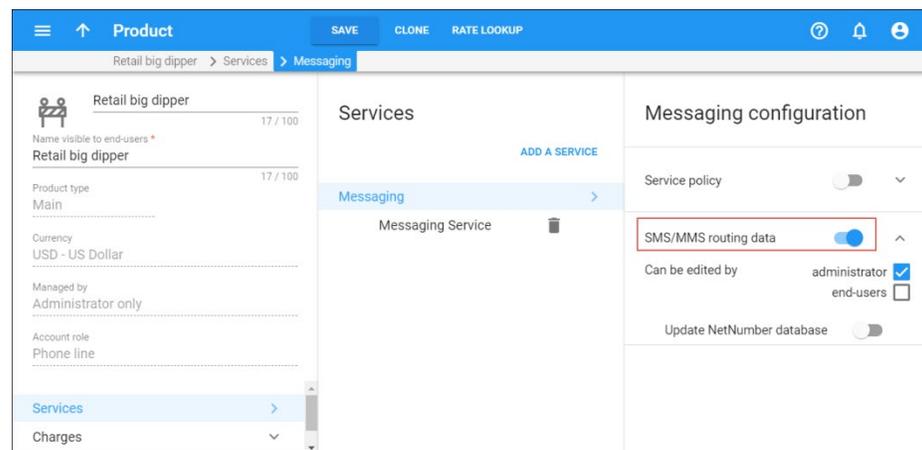
through Neustar Inc., but if necessary, PortaBilling® can collaborate with other number porting companies via additional plug-ins.

Integration with NetNumber

North American MVNOs use the NetNumber database for accurate SMS and MMS routing. This entails storing their DIDs in NetNumber and keeping them up-to-date.

PortaBilling® has the ability to notify NetNumber about the creation and termination of accounts. Upon receiving a notification, NetNumber updates the corresponding DID number configuration in their database. PortaBilling® also informs NetNumber if a number has been ported-in or ported-out from the system.

It is possible to configure whether to upload a particular account to the NetNumber database. If an end user is only subscribed to on-net services or simply does not need to receive SMS / MMS messages, by default, their account will not be uploaded to the routing database. For end users that require incoming SMS / MMS messages, a service provider purposely enables synchronization. It can be done for either the account's product or for the account itself.



NOTE: The synchronization with NetNumber must be enabled on the Configuration server. To use synchronization with the NetNumber database, configure options of the **NumberPorting_NetNumber** group on the Configuration server web interface.

The cooperation with NetNumber ensures easy and accurate routing for SMS and MMS in North America, thereby reducing messaging transmission costs for MVNOs.

Roaming identification in xDRs

With PortaBilling®, you can provide extensible information to your customers about the services they use while roaming and the charges they incur for this service usage. This information includes which country the user visits when using the service and the name of the carrier that provides this service.

Mobile carriers typically include the country (MCC) and carrier (MNC) codes in their CDRs to differentiate among services provided when users roam. Therefore, when an administrator obtains CDRs from a mobile carrier and imports them to PortaBilling®, the xDR Mediator extracts these codes and passes them to PortaBilling® for the called and / or calling party. For detailed description of xDR mediation utility, please refer to the **xDR Mediation – xDR Import** section.

PortaBilling® matches these codes against the name of the country visited and the carrier used and then adds this information to the xDRs produced for customers' accounts in the **More Info** dialog window on the xDR browser page.

xDR History for Customer 'John Doe'

Customer: John Doe
 From Date: 2016-11-22 00:50:51
 To Date: 2017-01-10 01:10:09
 Show xDRs For: Customer and Credit Accounts
 Service (s): Voice Calls

Total Transactions: 11
 Usage Charges: 15.22083 USD
 Usage Credits: 0.00000 USD
Total: 15.22083 USD

Voice Calls

?	View	More Info	Account	From	To	Country	Description	Date/Time	Charged Time, min:sec	Amount, USD	Hidden	Revert	
			4793003744@mob	4793003744	4748880990		Not Applicable	2016-12-02 08:53:03	0:08	0.00133			
			4793003744@mob	4793003744	4797059690		Not Applicable	2016-12-02 08:52:49	1:14	0.01233			
			4793003744@mob	4793003744	4790846394		Not Applicable	2016-12-02 08:50:40	1:27	0.01450			
			4793003744@mob	4793003744	4797059696		Not Applicable	2016-12-02 08:50:04	0:04	0.00067			
			4793003744@mob	4793003744	4799288193		Not Applicable	2016-12-02 08:50:04	0:52	0.00867			
Subtotal										3:45	0.03750		

Voice Calls, Roaming

?	View	More Info	Account	From	To	Country	Description	Date/Time	Charged Time, min:sec	Amount, USD	Hidden	Revert	
			4793003744@mob	4773790144	4793003744	(SPAIN, Vodafone)		2016-11-23 10:28:50	1:04	1.06667			
			4793003744@mob	4799274700	4793003744	(SPAIN, Vodafone)		2016-11-22 22:27:21	6:22	6.36667			
			4793003744@mob	4793003744	4799274700	(SPAIN, Vodafone)		2016-11-22 22:00:41	0:29	0.48333			
			4793003744@mob	4793003744	4791865123	(SPAIN, Vodafone)		2016-11-22 21:34:56	4:11	4.18333			
			4793003744@mob	4793003744	4748176284	(SPAIN, Vodafone)		2016-11-22 20:22:34	0:03	0.05000			
			4793003744@mob	4741291495	4793003744	(SPAIN, Vodafone)		2016-11-22 19:43:01	3:02	3.03333			
Subtotal										15:11	15.18333		

For customer and administrator convenience, PortaBilling® distinguishes between the xDRs produced for roaming services and those produced for home network service usage. Therefore, when a customer receives an invoice or browses their billing statistics on their self-care web interface, they see the roaming services (Voice calls, SMS and Internet services) and charges associated with them as a separate group.

By default, roaming xDRs are separated in an invoice template created in a layout designer. To adjust your custom invoice templates, please refer to the **PortaBilling® Templates Guide**.

Such enhanced representation provides a clear idea of charges incurred (both in the xDR browser and on the invoices) and helps to troubleshoot issues and solve disputes.

Twin card / data card services

You can introduce twin card and data card services with PortaBilling®. These are extra SIM cards that allow customers to use their service packages on multiple devices under the same subscription.

With the twin card, customers can use the service from two phones in exactly the same way as from a single phone. With data, SIM customers can share their data packages among multiple devices, e.g. between a phone and a tablet or a car. Any service used from the extra SIM is billed to the main account.

For twin card / data card services, use **beneficiary accounts**. Consider how it works:

John Doe is your mobile subscriber. His service bundle includes voice calls, SMS and 1GB of default data for a \$20 monthly subscription. Plus, he has several top-up options for his service wallet to add more data.

John orders an additional SIM card for his tablet and is charged \$5 for SIM card activation. Your administrator creates beneficiary account 12065558882 for his main account (12065552131) and assigns a SIM card to it. Now John can make calls and use the Internet from both his phone and his tablet.

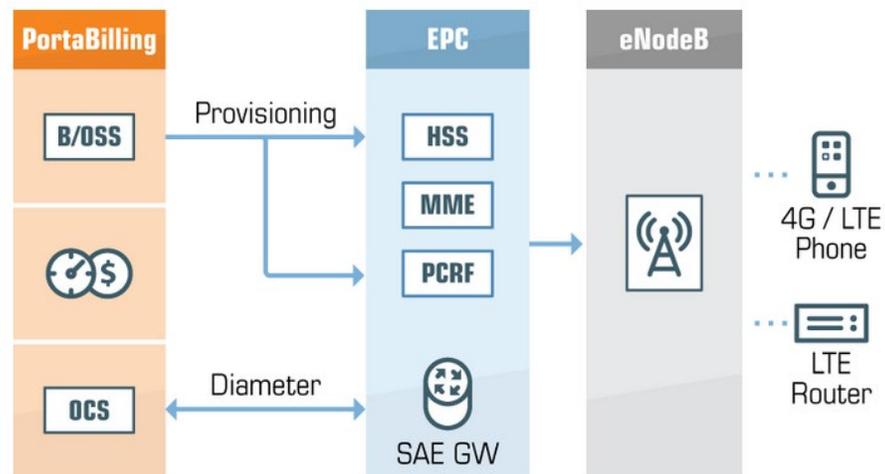
John uses up the 1GB of default data from his tablet and wants to purchase an additional 5GB. He enters his main account's self-care portal (12065552131) and tops up his service wallet for \$10. In his transaction history, John can see that his Internet sessions from the tablet are billed to his main account (12065552131).

7 ■ LTE service provisioning

LTE service provisioning and billing

LTE (Long Term Evolution) is becoming the new standard for building wireless networks. Numerous benefits in terms of radio technology (which are beyond the scope of this document) allow service providers to deploy the LTE network rapidly and at reduced cost. The large amount of bandwidth available to each consumer allows the provision of high-speed Internet combined with VoLTE and IPTV services. Standardized network infrastructure and interconnection between network elements (based on IMS) allow fast network expansion, easy replacement of some components and seamless interconnection with other operators.

PortaBilling can serve as the B/OSS system for LTE services, i.e. it manages customer data (including unified balance), service configuration details, product/rate tables, and so on for LTE just like for other services (VoIP or DSL Internet access). This allows the creation of a truly convergent environment and easy migration of existing customers (e.g. using WiMAX service) to the LTE network.



In terms of the network, PortaBilling performs the function of an OCS (Online Charging System). It communicates with SAE-GW via Diameter (Gy) for real-time session authorization and rating. The same Gy interface is used to disconnect a customer exceeding his credit limit or switch him to a different class of services when the usage quota is used up.

PortaBilling® provisions customer data to HSS (to permit centralized user authentication across the network) and PCRF (to allocate the correct policy / QoS parameters to the customer).

PortaBilling® can also provision static IP addresses to HSS. This functionality provides a variety of benefits, including the opportunity to access LTE devices from the Internet, provide remote support, set up

security cameras, set up a VPN connection and use the advanced features (e.g. port forwarding), among others.

NOTE: Static IP address provisioning to HSS is supported only for the Huawei SAE-GW and Aricent provisioning modules. Please contact the PortaOne® support team for assistance with how to provision a static IP address via a different manufacturer's gateways.

Precise real-time billing for period switch

In real-time billing for sessions that span several peak and off-peak periods, it is important to calculate the exact amount of traffic consumed during each time period.

To manage this task, PortaBilling® sends the exact time of the period switch to the gateway so that traffic consumption can be separately calculated for each period. Then the gateway sends these two figures to PortaBilling®. This allows providers of LTE, 3G, etc. to introduce precision in billing for separate time periods and preserve the continuity of the connection.

This is how it works:

A gateway sends recurring credit control requests to PortaBilling® in which the amount of traffic already consumed is reported. PortaBilling® sends the information about the traffic and time available to provide the service in the credit control answer. If PortaBilling® detects that the peak / off-peak switch will occur during the permitted time, then it includes the time of the switch in the credit control answer. The gateway then calculates traffic consumption separately for each period and sends these two numbers to PortaBilling®. PortaBilling® matches these numbers to peak and off-peak rates and creates two xDRs for charging the user accordingly.

NOTE: Your gateway must support the **Tariff-Time-Change** attribute.

To configure precise real-time billing for period switch, specify *Yes* for the **tariff-time-change** option in the Internet Access policy. For more details and configuration guidance, see the **How To ...** handbook of the *LTE Provisioning* section).

Rating groups

ISPs can differentiate among types of Internet traffic. This enables them to bill customers differently, based on the content they access (e.g. provide free access to your hotline portal while charging users for watching movies online). Such content-based billing is implemented via **rating groups**. Their usage provides extra space for special offers and

promotions and can be used to influence customers' usage patterns, thereby increasing company revenues.

Let's consider the following example. A service provider offers Internet access free of charge to the following addresses:

- Service provider's hotline website
- Social networks, and
- Service provider's SIP server (for making SIP calls).

Everything else is billed using a standard price. To configure this, the service provider defines the rating groups on the network side (gateway, PCRF, etc.) and assigns a unique index to each of them. For example:

- Service provider's hotline website – RG101
- Social networks – RG102
- Service provider's SIP server – RG103
- Regular web browsing – RG104.

When an end user visits a website, the gateway sends the *rating-group* attribute that defines this type of Internet traffic to PortaBilling®. PortaBilling® matches the rating-group attribute with a corresponding rate in the customer's tariff and applies the charge. Note that the NETACCESS rate must also be present in the tariff to authorize a user and charge for any other rating group.

For detailed information about how to configure content-based billing, refer to the [Unified PortaSwitch Handbook Collection](#).

Multiple rating group authorization

A customer can use different content within the same session (e.g. open a Facebook page and watch a YouTube video). In this case, the node sends two rating groups and requests the amount of Internet quota for each group in the authorization request.

PortaBilling® authorizes all rating groups provided within the same authorization request. To make this happen, it supports the Multiple-Services-Credit-Control (MSCC) AVP for Diameter(Gy). When P-GW sends the MSCC group in the Credit Control Request - Initiate (CCR-I) message to PortaBilling®, PortaBilling® authorizes rating groups within MSCC AVPs one by one, starting from the most expensive.

To illustrate how this works, consider the following example:
Let's say you introduce access to different Internet content types via add-ons.

You have configured the P-GW to send rating group 101 for YouTube watching and rating group 102 for online games. Other types of Internet access (e.g. regular web browsing, emails, etc.) are authorized using the NETACCESS rate in a customer tariff. The P-GW requests 100MB for each type of Internet content in the authorization request.

A user product configuration includes a default 500MB quota that can be used for any online activity.

John Doe has signed up for 1GB of YouTube watching for \$5 / month. He has already used up 950 MB and has only 50MB available.

When on the next day he connects to the Internet to check his emails and watch YouTube, PortaBilling® authorizes his session as follows:

- For rating group 101 PortaBilling® first checks the amount of quota available for YouTube. Since the user has only 50 MB left, PortaBilling® allocates this amount, plus allocates 50 MB from the default 500MB quota.
- PortaBilling® authorizes the user for email check using the NETACCESS rate and allocates 100MB from the remaining 450MB (500 MB – 50MB = 450MB) of default quota.
- When 100MB allocated for YouTube is consumed, the P-GW sends another authorization request for rating group 101. PortaBilling® authorizes it and allocates another 100MB from the user's default quota.

Thus, users are protected from overdraft and PortaBilling® charges correctly and precisely for their Internet usage.

Billing for the LTE service in roaming

To stay competitive in the mobile market, wireless operators provide their subscribers with roaming Internet service when their subscribers visit other regions or countries. Wireless operators can bill their subscribers differently – according to the country and mobile network within which the Internet is actually used. The country and network are identified by the MCC / MNC codes.

NOTE: The ASN-GW must send the 3GPP-SGSN-MCC-MNC attribute in Diameter (Gy) requests.

Consider the following example:

MVNO GoMobile provides home network Internet access for \$0.1/MB. For roaming Internet, GoMobile charges \$0.5/MB in Australia and \$0.7/MB in the UK.

John Doe is going to travel abroad so he tops up his available funds by \$100. In Australia he uses 20 MB of roaming Internet service for \$10 ($\$0.5/\text{MB} * 20 \text{ MB} = \10). And from Australia, John visits the UK where he uses 35 MB of roaming Internet service and is charged \$24.5 ($\$0.7/\text{MB} * 35 \text{ MB} = \24.5).

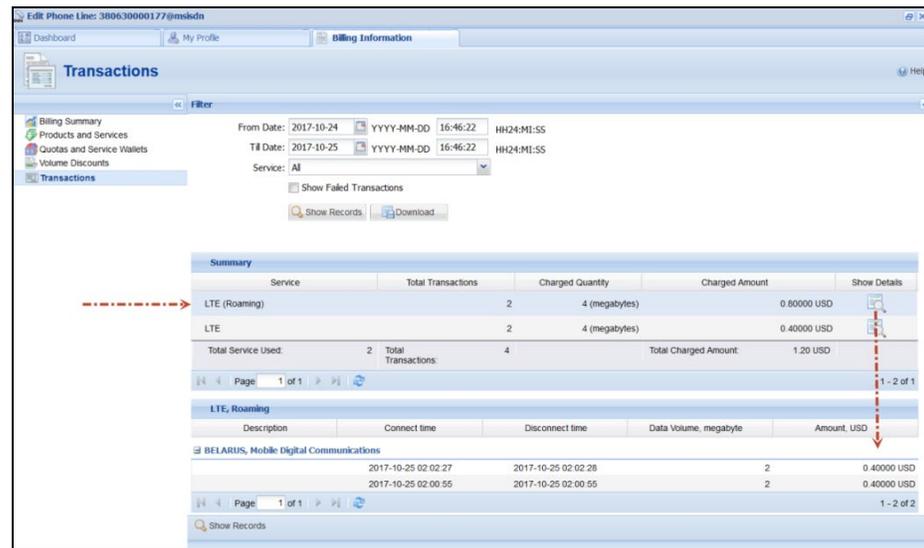
At the end of the trip John still has \$65.5 left ($\$100 - (\$10 + \$24.5) = \65.5). When John returns to his home network and begins browsing the Internet, he is once again charged \$0.1/MB.

To configure the billing for roaming Internet service based on the country and network, a wireless operator must:

- enable the **RateMatchWithE212** option on the Configuration server, and
- create a customer tariff in the E.212 format.
- specify the home PLMN (Public Land Mobile Network) in the Internet access policy (e.g. `home-plmn-e212-match = 310`).

When the home PLMN is configured, the roaming xDRs are grouped separately on invoices and in the xDR browser / Transactions page. Thus, when subscribers log into their self-care web interface, they can clearly differentiate Internet services consumed within their home network from those used in roaming. The detailed view provides information about the roaming country and charges.

For instance, John Doe wants to review the costs for Internet usage in roaming mode. He browses the billing statistics on his self-care web interface, finds the records for roaming services grouped together and reviews the roaming charges.



It is possible to configure the billing for national roaming Internet. Please contact the PortaOne® support team for more details.

By using this solution, wireless operators can bill subscribers differently based on the MCC / MNC and thereby improve their cash-flow.

Individual hotlining

Hotlining is a convenient way to redirect end users who subscribe to WiMAX, LTE or 3GPP services to a defined website. There they can quickly find information about how to subscribe or prolong their Internet service if their traffic quota or balance are insufficient for using the Internet. Different services can be provided (e.g. for residential and airport hotspots; resellers, etc.) and a hotlining link can be configured for each of them using individual Internet access policies.

To configure hotline redirection, create an Internet access policy and assign it either to a product or directly to an account. Input the main configuration parameters on the **Hotline** tab:

- **When account out of funds** – Select this option to redirect end users if their traffic quota or balance are insufficient for accessing the Internet.
- **When account suspended or blocked** – Select this option to redirect end users whose accounts are blocked, suspended or expired.
- **Captive portal URL** – Specify the website in the URL format where end users are redirected to.

The screenshot shows the 'Service policy' configuration page in the PortaBilling interface. The 'Hotlining' tab is active, and the following options are checked:

- If account is out of funds
- If account is suspended or blocked

The 'Captive portal URL' is set to `https://mycompany.com`.

NOTE: Hotlining is a feature that relies on a third-party network solution. Make sure that the networking platform (NAS, ASN, etc.) supports hotlining and is properly configured.



Note that services provisioned via the Diameter protocol require the billing engine's DIAMETER module to be enabled on the Configuration server. On that same page it is possible to globally set the hotlining link for all PortaBilling® environments.

PortaBilling currently supports the Diameter Credit-Control Application (**DCCA**) with additional AVPs defined in **3GPP TS 32.299** (version 12.0.0 Release 12).

Configuration for unregistered accounts

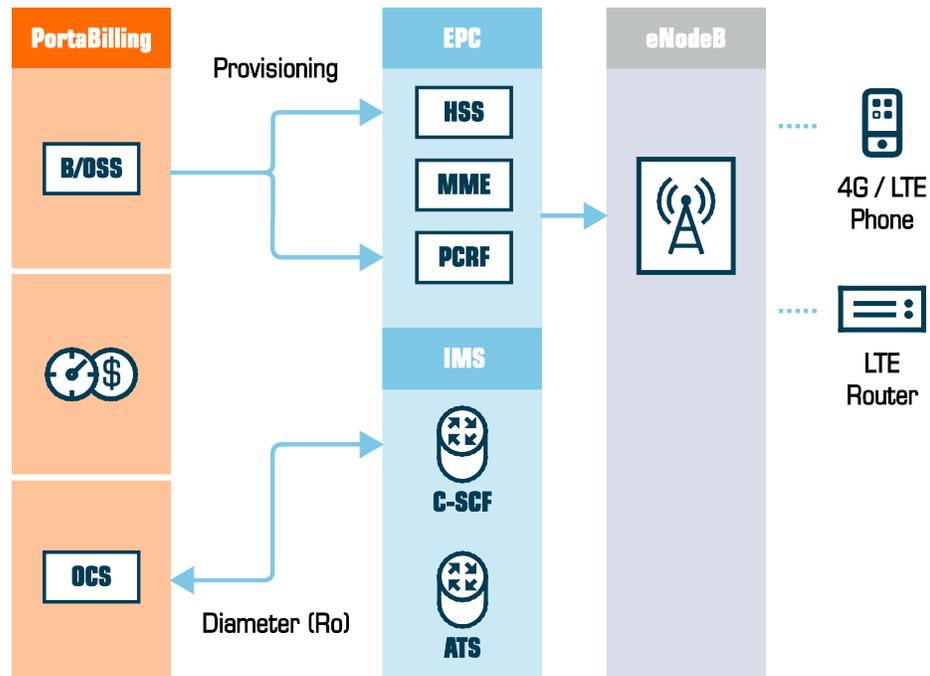
It is also possible to configure hotlining for end users who are not registered in PortaBilling® instead of just rejecting the session. To do this, perform the following steps:

1. Create an account with the UNKNOWN@msisdn account ID.
2. Create an Internet access policy with configured hotline redirection.
3. Assign the Internet access policy to the account.

If PortaBilling® does not find the account that attempts to access the Internet service, the system uses the UNKNOWN@msisdn account configuration and redirects the user to the defined webpage.

VoLTE

Voice over LTE (VoLTE) is the solution for providing voice traffic on an LTE network. It allows you to utilize all the benefits of LTE: quick user connectivity, better service quality and higher performance, etc. It also helps to ensure a better LTE experience for your customers by allowing them to use their favorite data and voice applications simultaneously.



To deploy VoLTE services with PortaBilling® as the online-charging system (OCS), PortaBilling® supports the Diameter (Ro) interface. Thus, PortaBilling® communicates with IMS via Diameter (Ro) for real-time user authorization and charging. When a user runs out of funds or exceeds their credit limit, PortaBilling® instructs the node to disconnect the user, thereby preventing revenue leakage.

PortaBilling supports processing of calls to ported numbers in IMS networks in the same way as in VoIP networks. See more in the [Local number portability](#) section.

By performing the functions of an OCS, PortaBilling® manages the customer balance and service configuration (service policies, quotas, etc.) that are provisioned to other components on the LTE network.

Supported features

In conclusion, the list of features available in PortaBilling® and provisioned via the Diameter protocol can be divided into the following categories:

- General features for both Diameter (Gy) and (Ro) interfaces,
- Features provisioned via Diameter (Gy) and
- Features provisioned via Diameter (Ro).

General features

- Online Centralized Session charging with units reservation (SCUR) and Credit-control application (DCCA).
- Charging based on the time of use (peak / off peak).

Features provisioned via Diameter (Gy)

- Traffic categorization into rating groups.
- Hotlining.
- Server-initiated credit re-authorization (RAR) to lift service restrictions from an account upon account top-up. It is triggered by ESPF events.
- Additional AVPs (Volume-Quota-Threshold, Time-Quota-Threshold, Quota-Consumption-Time, Quota-Holding-Time, Multiple-Services-Credit-Control).

Features provisioned via Diameter (Ro):

- Online Centralized Event (SMS) charging with units reservation (ECUR).
- Voice calls charging (e.g. outgoing off-net / on net calls, incoming calls).

With further PortaBilling® interconnection with the mobile network, this list will be extended with other features in upcoming releases.

Please find Diameter supported RFCs and PortaBilling® communication diagrams in the **Appendices** section.

Verified devices

PortaBilling® is integrated with a range of devices that enable wireless operators to provide full-scale LTE services. These devices have been tested and verified in PortaOne’s laboratory as interoperating with PortaBilling® successfully.

All devices can be divided based on the Diameter interface they use for communication with PortaBilling® (Gy or Ro) and the features they support:

Devices operating in Diameter Gy

Device	Supported Features
Huawei UGW9811	<ul style="list-style-type: none"> • Traffic categorization into rating groups • Credit Control for Multiple Services per session • Hotlining • Server-Initiated Credit Re-Authorization (RAR) • Online Centralized Session charging with units reservation SCUR, DCCA • Real-time billing for period-switch
Cisco ASR 5000	<ul style="list-style-type: none"> • Traffic categorization into rating groups • Credit Control for Multiple Services per session • Hotlining • Server-Initiated Credit Re-Authorization (RAR) • Online Centralized Session charging with units reservation SCUR, DCCA • Real-time billing for period-switch

ZTE ZXUN xGW	<ul style="list-style-type: none"> • Traffic categorization into rating groups • Credit Control for Multiple Services per session • Hotlining • Server-Initiated Credit Re-Authorization (RAR) • Online Centralized Session charging with units reservation SCUR, DCCA
Adax EPC	<ul style="list-style-type: none"> • Traffic categorization into rating groups • Credit Control for Multiple Services per session • Server-Initiated Credit Re-Authorization (RAR) • Online Centralized Session charging with units reservation SCUR, DCCA
YateUCN EPC	<ul style="list-style-type: none"> • Traffic categorization into rating groups • Credit Control for Multiple Services per session • Hotlining • Server-Initiated Credit Re-Authorization (RAR)
Protei DPI	<ul style="list-style-type: none"> • Traffic categorization into rating groups • Credit Control for Multiple Services per session • Hotlining • Server-Initiated Credit Re-Authorization (RAR) • Online Centralized Session charging with units reservation SCUR, DCCA

Devices operating in Diameter Ro

Device	Supported Features
BroadForward BFX Interface Gateway	<ul style="list-style-type: none"> • Online Centralized Event (SMS) charging with units reservation (ECUR).
Huawei ATS ATS9900	<ul style="list-style-type: none"> • Online Centralized Event (SMS) charging with units reservation (ECUR). • Voice calls charging (e.g. outgoing off-net / on net calls, incoming calls).
Protei CAMEL Gateway	<ul style="list-style-type: none"> • Online Centralized Event (SMS) charging with units reservation (ECUR). • Voice calls charging (e.g. outgoing off-net / on net calls, incoming calls).

Protei SMSC	<ul style="list-style-type: none">• Online Centralized Event (SMS) charging with immediate event charging (IEC)
YateUCN	<ul style="list-style-type: none">• Online Centralized Event (SMS) charging with units reservation (ECUR).

PortaBilling® is also capable of communicating with other devices and remote gateways using the Diameter protocol. If you wish to verify the interoperability between your device and PortaBilling®, please contact the PortaOne support team.

Supported HSS / HLR

PortaBilling® is integrated with the following HSS / HLR for SIM card provisioning:

- Adax HSS
- Aricent HSS
- Huawei HSS
- Greenis HSS
- NetNumber Titan HSS
- Protei HLR / HSS
- Yate HSS/HLR
- ZTE HSS

Please consult the [PortaSwitch® Interoperability Guide](#) for a detailed description of PortaBilling® integration with the listed HSS / HLR.

8 ■ VoIP services

Service policies

The telecommunications industry is growing rapidly with new technologies and devices being introduced into the market all the time. Networks have become more flexible and customers have become more demanding about the services that they subscribe to. Under these conditions, it is essential to not only continually offer new services and products and keep up with market demand, but also to do this in the most qualitative and flexible manner. This is what PortaBilling® enables you to do!

The **Service policies** feature allows you to fine-tune your services based on your network peculiarities, vendors' opportunities and customers' demands. It facilitates the configuration of static options for multiple accounts (so it is not necessary to configure values for each account separately), thus establishing common policies for groups of accounts. It also allows you to separate the technical configuration of specific options (usually made by technical staff) from account management.

When a new policy is created, it does not have any attributes defined (all available attributes are shown in grey). To define an attribute, specify its value.

Service policies can be statically assigned at various levels: account, product, connection and authentication (call handling rule). In addition, it is possible to apply service policies depending upon the UA type (dynamically matched policy).

Dynamically matched policies

The service policy has the ability to match dynamically when the **Match Pattern** is specified. For the calling party, PortaSIP extracts the *User-Agent* header from the incoming INVITE request (e.g. "Linksys/SPA941-5.1.8") and matches it against all of the dynamic policies. If more than one policy matches, the one with the highest **Match Priority** is used.

For the called party, the procedure is quite similar. The only difference is that the *User-Agent* header of a called UA is taken from its registration information.

Policy precedence for VoIP services



Note that service policies apply for both call parties (caller and callee) individually. The final set of attributes for each party is derived from service policies applied to different stages (a service policy matched on initial request is also considered).

To illustrate how this works in practice, consider the following example: Let's say John Doe makes a call to number 12125554521. You want to ensure that the call passes via the G.711 codec so you configure the service policy for John's account and define G.711 as the first codec for the `codec_order_list` option.

When John calls, the service policy assigned to his account is applied and this reorders the codecs.

The call is routed to the ABC vendor who requires the user identity of the user making the call. Therefore, the service policy assigned to the ABC vendor connection instructs PortaSIP® to send the user identity information in the INVITE request.

What happens if there is a contradiction in service policies? For example, if a service policy matched upon initial request has the *no_voice_rejects* attribute set to Yes, while the service policy assigned to a called party has its attribute value set to No?

Several simultaneously applied service policies function according to the precedence of defined attributes that are exact matches:

- If there is a policy that corresponds with the SIP endpoint name of the called party, then this policy's attributes are applied.
- If the called party has a service policy assigned via a product, then that overrides the one matched upon initial request and its attributes are applied. However, if an administrator overrides a service policy for an account, this policy takes the highest priority.
- The same attribute can be defined in the service policy assigned to the call handling rule and the one assigned to the incoming connection. In this case, the service policy assigned to the connection takes higher precedence and its attribute value is applied.

Built-in policy attributes

There is an attribute-specific prevalence among service policy attributes assigned to both caller and called accounts:

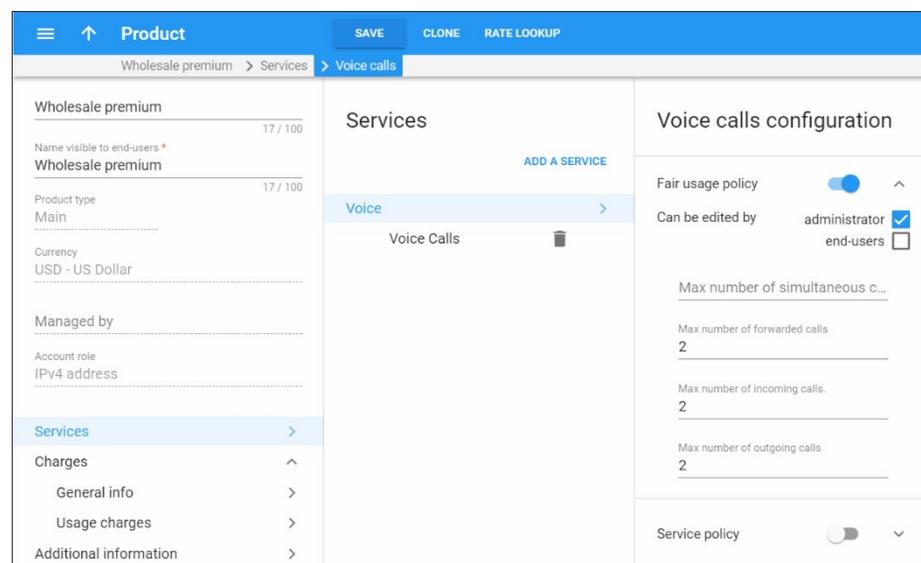
- The *codec_order_list* attribute is taken from the caller's account (called account's attribute is ignored).
- The "header" attributes (*out_hdr_pai*, *out_hdr_rpid*, *out_hdr_history*, *out_hdr_diversion*) defined for a called account take precedence over the ones defined for a caller account.
- The *keep_alive_interval* attributes are individually defined for both the caller and called parties, so there is no dominance between them. The value for the calling party is taken from the service

policy assigned to that account. The value for a called party is taken from the service policy assigned to the connection or the one that's assigned to the account if the called party is one of your accounts.

Limiting per-account simultaneous calls on a product

Fair usage policy is very important for every ITSP. Let's assume that you have a product that supplies your residential customers with an unlimited number of minutes for making free calls to the USA. If you don't want someone to use those free minutes via their PBX by making many simultaneous calls, you can set the limit for two simultaneous calls.

Fair usage policy is configured on the Product's **Service Configuration** tab, where this functionality can be enabled and the limits defined for simultaneous, incoming, outgoing and forwarded calls.



The screenshot displays the 'Product' configuration page for 'Wholesale premium'. The 'Services' section is active, showing a 'Voice Calls' service. The 'Voice calls configuration' panel on the right is expanded, showing the following settings:

- Fair usage policy: (enabled)
- Can be edited by: administrator and end-users
- Max number of simultaneous c...: 2
- Max number of forwarded calls: 2
- Max number of incoming calls: 2
- Max number of outgoing calls: 2
- Service policy: (disabled)

This functionality can be useful for the following situations:

- You have a SmartCall product that supplies your customers with free domestic calls. There is a new customer who buys this product from you. You supply them with credentials to use the services. To ensure that services are for personal use only and that nobody can use them with an office PBX (to send hundreds of calls from their call center free of charge), the administrator can set the **Maximum Number of Simultaneous Calls** for the SmartCall product to 2.
- The administrator sets the **Maximum Number of Outgoing Calls** to 100 for the SmartCall product. This product is assigned

to the customer who has his own call center. If the customer needs more simultaneous outgoing calls he can buy another product that offers that.

- The administrator enables call forwarding functionality and sets the **Maximum Number of Forwarded Calls** to 2 for the SmartCall product. This allows each user to forward 2 calls coming simultaneously to their SIP phone to their mobile and landline numbers when they are out of the office.

PortaBilling® only limits billable calls – calls for which PortaBilling® produces xDRs and applies charges to an account. These are transferred calls (charged with the tariff matched by the TRANSFER access code), redirected calls (e.g. calls that arrive to an auto attendant from external networks) and calls forwarded to off-net destinations (charged with the tariff matched by the FOLLOWME access code). If a final forwarding destination is a SIP account, the limit on forwarded calls is not applied.

For more detailed information on limiting forwarded calls consult *Limiting calls forwarded to off-net destinations* section in the **PortaSIP® Administrator Guide**.

This functionality can help you prevent fraud and service abuse of the products that provide your customers with free calls.

Limits on calls in progress

Sometimes a customer's SIP gateway (e.g. an auto-dialer) produces many call attempts when trying to reach a destination that does not answer, and by doing so, the system becomes overloaded.

Another example is “cross-forwarding” in which forwarding is configured for the same user, both in the SIP network (e.g. from an extension to a mobile phone) and in the mobile network (e.g. from the mobile phone to an extension). As a result, the call “loops” and is processed as separate call attempts to the mobile network.

These numerous call attempts consume system resources for authorization and routing, and consequently slow down the system and affect normal call processing.

Therefore, to increase system stability and prevent its overload, both connected and calls in progress are considered when limiting the number of simultaneous calls.

As soon as a limit is set on simultaneous calls for a given customer, customer site or account (as part of product configuration), PortaBilling®

requests notifications from PortaSIP® about this customer's calls in progress.

This is so that if excessive call attempts occur, the limit on simultaneous calls will quickly be reached and subsequent call attempts will be rejected. The administrator can also track calls that are in progress on the **Active Sessions** page.

The ability to limit connecting calls serves as an additional tool against DoS attacks while improving overall system security.

Customer Premises Equipment (CPE) provisioning and inventory

If you provide your VoIP customers with CPE equipment like IP phones, you know how laborious and yet important the task of performing initial configuration is. If the equipment is not configured properly, it will not work after being delivered to the customer. Or, even if it works initially, problems will arise if you need to change the IP address of the SIP server. How can you reconfigure thousands of devices that are already on the customer's premises? There are two ways to manage the device configuration.

Manual provisioning

The administrator must login to the device provisioning interface (typically HTTP) and change the required parameters. There are several drawbacks to this method:

- The IP phone must be connected to the Internet when the administrator is performing this operation.
- The administrator must know the device's IP address.
- The IP phone must be on the same LAN as the administrator, or on a public IP address (if the device is behind a NAT/firewall, the administrator will not be able to access it).

Due to these reasons, and since every device must be provisioned individually, this method is acceptable for a testing environment or small-scale service deployment, but totally inappropriate for ITSPs with thousands of IP phones around the world.

Auto-provisioning

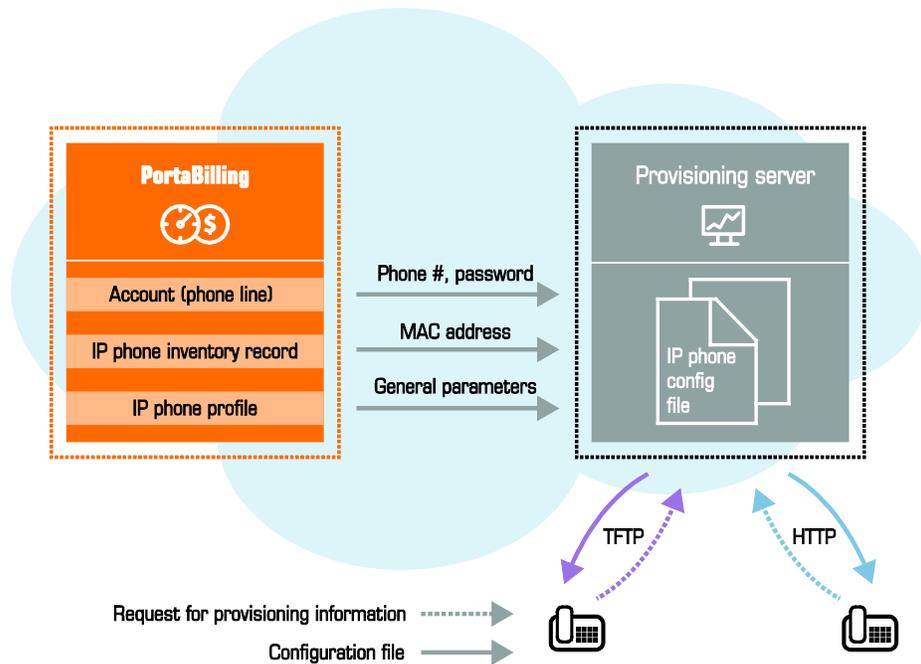
This approach is a fundamentally different one. Instead of attempting to contact an IP phone and change its parameters (pop method), the initiative is transferred to the IP phone itself. The device will periodically go to the provisioning server and fetch its configuration file.

IP phone provisioning

When you use auto-provisioning for an IP phone, instead of entering the same values for codec, server address, and so on into each of a thousand user agents, you can simply create a profile which describes all these parameters. Then PortaBilling® can automatically create a configuration file for the SIP phone and place it on the provisioning server.

The only configuration setting which is required on the IP phone side is the address of the provisioning server, i.e. where it should send a request for its configuration file. When the IP phone connects to the Internet, it will retrieve a specific configuration file for its MAC address from the TFTP or HTTP server and adjust its internal configuration.

If you decide later to change the address of the SIP server, you need only update it once in the profile, and new configuration files will be built for all user agents. Each user agent will then retrieve this file the next time it goes online.



The config file is specific to each user agent, as it contains information such as username and password; thus the user agent must retrieve its own designated config file. The following are defined in the billing configuration:

- The CPE profile, so that the system knows which generic properties (e.g. preferred codec) to place in the configuration file.
- An entry about the specific IP phone in the CPE inventory (including the device’s MAC address), with a specific profile assigned to it.

- The IP phone (or, in the case of a multi-line device, a port on the phone) is assigned to a specific account in the billing.



To speed up configuration file generation (e.g. when auto-provisioning large numbers of IP phones (1000 and more)), you can make a parallel processing mode available for this task. To do this, enable the **UA_Profiles.Enable_Parallel_Generation** option on the Configuration server.



Auto-provisioning will only work if your IP phone knows the address of your provisioning server. If you buy IP phones retail, you will probably have to change the address of the provisioning server on every phone manually. However, if you place a large enough order with a specific vendor, these settings can be pre-configured by him, so that you may deliver an IP phone directly to the end user without even unwrapping it.

Some IP phones can be customized by end users. Thus, end users may wish to configure particular options (e.g. soft keys) according to their preferences. Then, you can perform more specific configurations of IP phone auto-provisioning. In the CPE profiles, you can specify options to be auto-provisioned by PortaSwitch® (e.g. the addresses of SIP servers, codecs, passwords, etc.) and exclude options that end users can configure themselves.

For example, when creating a CPE profile for Audio Codes 405HD, remove such personal settings as programmable keys from the list of options to be auto-provisioned. Then, if end-users configure the programmable keys for their IP phones, those changes remain after the IP phone configuration update.

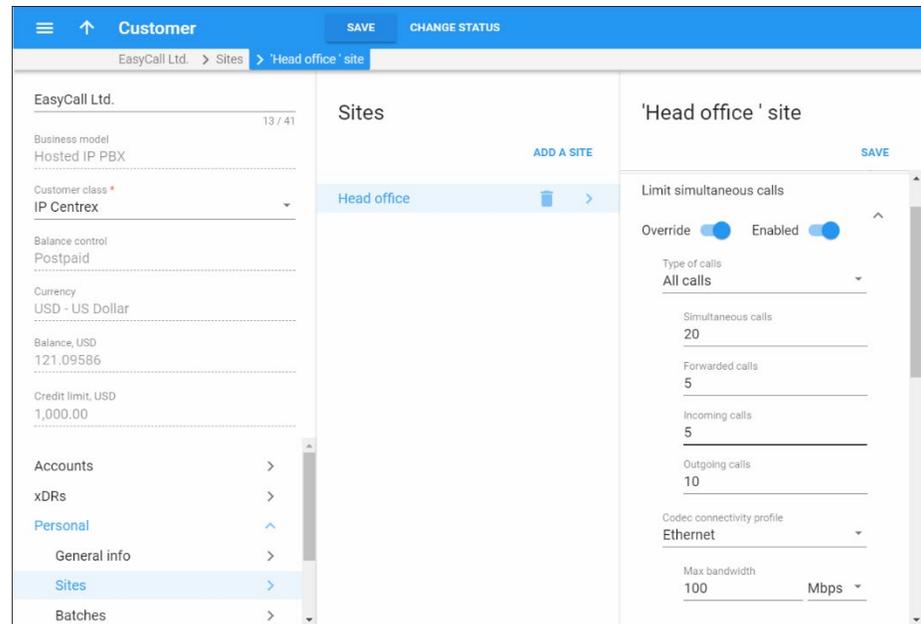
CPE Inventory

The CPE inventory allows you to keep track of IP devices (SIP phones or adaptors) which are distributed to your customers. The MAC address parameter is essential for every IP phone which is to be automatically provisioned, and so a corresponding entry must be created in the CPE inventory.

Bandwidth-aware limitation on concurrent calls permitted

In order to guarantee a constant adequate quality of service in a situation where a customer's Internet connection has limited bandwidth, it is important to consider not only the total number of already established concurrent calls, but also how much bandwidth each one consumes. (The actual amount of bandwidth depends on the codec used and other

parameters.) A new addition to the call control module allows the administrator to specify limitations in terms of the available bandwidth, so that PortaSwitch can then correctly determine whether a specific customer's site can place (or receive) additional calls.



Active call monitoring

PortaBilling® can collect information about calls in progress from the PortaSIP® server, Cisco/Quantum gateways and other VoIP nodes. The only thing required on the VoIP node side is the capability to send Start accounting records, in addition to Stop records.

This enables several options:

- You may see a list of calls in progress on the PortaBilling® web interface.
- You may forcibly disconnect a specific call (if the call is routed via PortaSIP®).
- You may specify the maximum allowed number of simultaneous calls for each customer. PortaBilling® will reject any call attempts above the allowed number of concurrent calls.

Using different protocols (SIP and H323)

The billing itself is independent from the actual implementation details of those protocols used to carry out the call. PortaBilling® communicates with the nodes on your VoIP network (gateways, gatekeepers, proxies) using the RADIUS protocol and the VoIP VSA extensions for it. Thus, from the point of view of PortaBilling®, RADIUS requests related to SIP

calls are nearly identical to those related to H323 calls (there are only a few different attributes).

PortaBilling® is able to process both types of requests, so no special configuration in billing is required to support SIP or H323 calls. You only need to add your gateway to the system as a trusted node, and then both H323 and SIP calls reported from that gateway will be processed properly.

NOTE: From the billing point of view, a SIP server is not different from any other gateway - it also issues authentication or authorization requests, and sends accounting information after the call is completed.

A node for the SIP server should therefore be created in your system, so that it can communicate properly with the billing. Also, products that are to be allowed to use SIP services should have a rating entry which includes this node.

9. Fraud protection tools

Tools for prevention of VoIP fraud

PortaSwitch® supports a set of tools to prevent VoIP fraud, i.e. a situation where an attacker uses credentials stolen from a customer's IP PBX or end-point and sends unauthorized voice traffic to your network.

There are two main approaches for distinguishing between the activity of a legitimate user and a hacker:

- Deviation in usage patterns.
- Deviation in the location from which the service is used (geo-IP).

The current version of the fraud protection module in PortaSwitch uses the location deviation method. This enables prevention of the majority of attacks without allowing a single fraudulent call to pass through, so there is zero loss for the service provider (or the customer). We plan to add real-time analysis of usage patterns in the next releases, to allow even better control and proactive prevention of those rare scenarios where location analysis may not yield sufficient results.

Geo-IP fraud prevention

If a hacker obtains the valid credentials of one of your customers, they can then send a call from their network using this username and password and even the identification of the SIP phone type. The only piece of information that cannot be easily changed by a hacker is their actual IP address. Although it is not always possible to determine a user's exact location (e.g. street address) from their IP address, information regarding the country and the ISP that owns the IP address (and leases it to the end user) can be determined in the overwhelming majority of cases by using the database of IP address assignments (geo-IP database).

The key element in geo-IP fraud prevention is the assumption that, under normal circumstances, the majority of users will use your service from the same country (and the same region in that country). For instance, a VoIP user in Barcelona, Spain, connects to the Internet via DSL provided by Telefonica. Although the user's IP address is assigned dynamically by the ISP and may change, the "location" of that address will always point back to Spain. Even if the user changes the ISP to Orange, their IP address will still refer to Spain. Usually, larger customers who use SIP trunking services are the primary target for hackers, since they have higher credit limits and are authorized to send larger numbers of simultaneous calls. By using the credentials of such a customer, a hacker can send a significant amount of traffic in a relatively short time period. Fortunately, geo-IP fraud prevention is especially accurate for this type of customers. A residential user may take his SIP phone with him to a different country

while on vacation, but it is very unlikely that a company will frequently move their office (and IP PBX) to a different country.

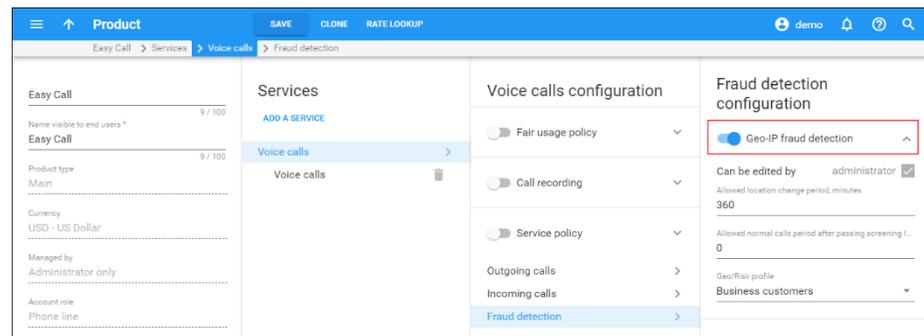
When launching the service, the service provider creates one or more **Geo / Risk Profiles**. Each profile divides the list of all the countries in the world into three zones:

- Trustful – A country (or countries) where users intend to use the service. Service usage is allowed without restrictions.
- Cautious – Countries where it would be unusual (but still possible on a relatively low number of calls) for a customer to use the service without screening. Then any attempt to make an outgoing call from a country listed here will be screened, where he or she must provide additional credentials to prove that this is indeed a legitimate user.

NOTE: The number of suspicious calls that can be made without screening is configured for each Geo / Risk profile. The default value is 5, you can change it to 3 or 10.

- Paranoid – Any usage attempt from these countries will be treated as a potential security breach and immediately screened. The customer must provide additional credentials to prove that this is indeed a legitimate user.

Naturally, the actual selection will depend on the area where the service provider sells the service and what type of service it is. For instance, an ITSP selling business SIP trunking services in London may define the United Kingdom as the “trustful” service location, and everywhere else as cautious, since all of their customers are actually based in the UK. Another ITSP, selling residential VoIP calling via a dialing app on a smart phone will define the United Kingdom, France and Spain as trustful locations (since they actively advertise their product in these countries and most of their customers are there). In this case, the majority of countries in the world will be listed as “cautious.” Finally, those countries from which the service provider sees an increased amount of hacking attempts will be declared “paranoid.”



Next, Geo-IP fraud detection can be enabled in individual products, thereby allowing IP verification to be performed for all accounts, using this product. This allows you to apply Geo-IP verification to business and residential VoIP products, and skip it for other products (for backward compatibility or simply because it does not fit the business model, for instance for a service similar to Skype, where users can register and use the service anywhere in the world).

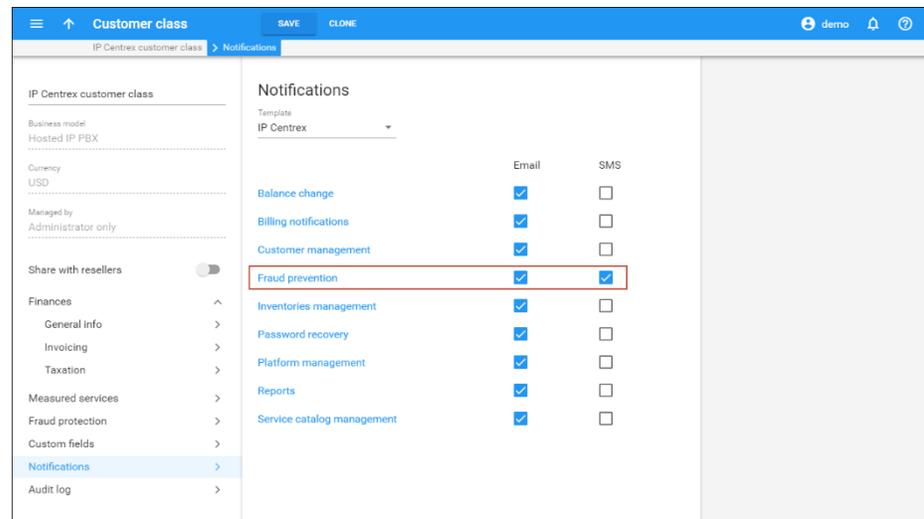
Alerts about suspicious activities

Time is critically important when a fraudulent incident occurs. The sooner one acts, the better chance they have to mitigate possible losses.

PortaSwitch® provides notifications to alert customers whenever the following events take place:

- A call made from a customer’s account originates from a location known for high fraud potential and is therefore redirected to the screening IVR. The account’s status is changed to **Screened**.
- A call made from a customer’s account originates from a location known for high fraud potential and is therefore immediately rejected. The account’s status is changed to **Quarantined**.

To automatically notify your customers about these events, the **Fraud prevention** group of notifications must be enabled in the customer class.



When customers are promptly notified about suspicious activities, they are able to detect fraud early and take necessary protective measures to avoid possible losses.

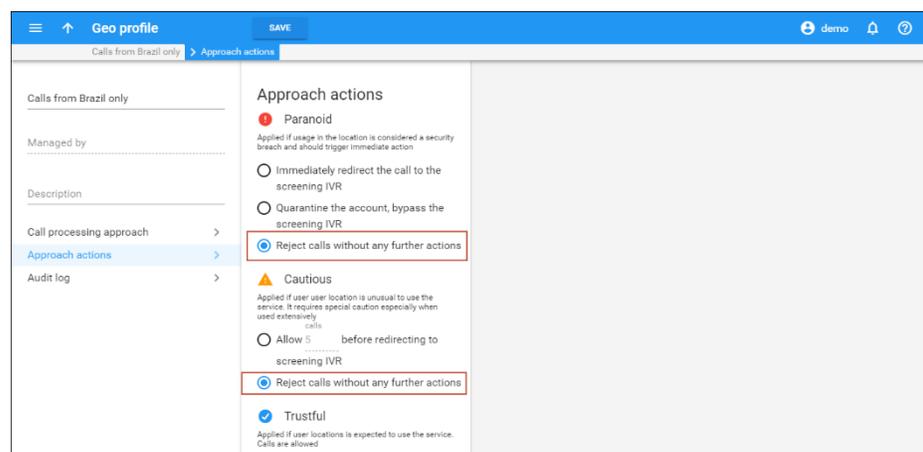
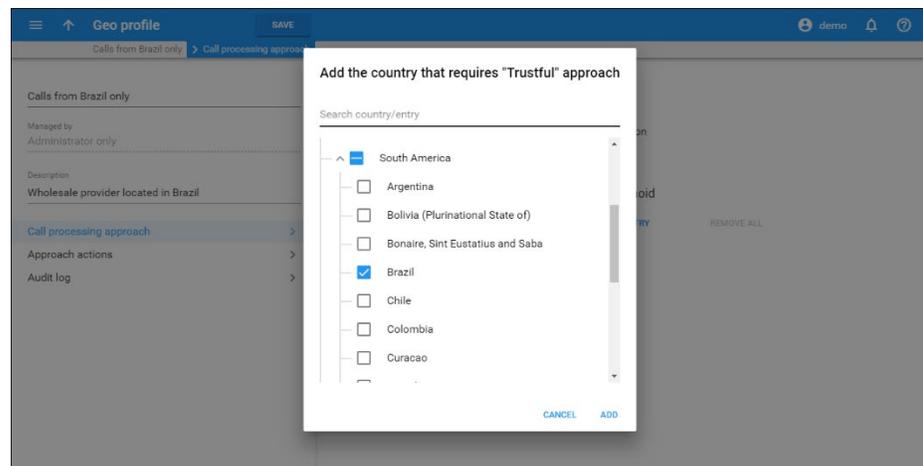
Country blacklist

If the service provider wants to block calls originating from countries known for high fraud potential and the screening of calls by means of IVR is unacceptable, the Country blacklist functionality is the one to use.

This functionality immediately drops calls that come from blacklisted countries. While fraudulent calls are rejected, calls originating from trusted countries come through without restrictions.

Consider the following example. A customer who is a wholesale provider is located in Brazil. He wants to allow calls only originating from Brazil and deny calls originating from any other country. As many calls are sent simultaneously, screening calls by means of IVR is not optimal.

For this customer, an administrator configures the Country blacklist to include all countries except Brazil and instructs PortaBilling to deny calls from these countries by default.



After a while, a Chinese hacker hacks this customer's account and attempts to send calls to expensive destinations. PortaSwitch® detects that the account is attempting to send calls from a Chinese IP address and since China is present in the Country blacklist, these fraudulent calls are rejected, although calls from Brazil come through without any restrictions.

With Country blacklist functionality, service providers can offer more enhanced fraud protection to their customers – including the ability to consistently drop calls that originate from countries with a high fraud risk and to continue sending calls from trusted countries only.

Account attributes

The default setting for an account is “stationary,” meaning that the system expects the account to always be used in the same location. A change of country (location) is immediately considered a potential security breach. The alternative is to mark some accounts as “roaming,” in which case, the system allows for a greater number of location changes. This can be useful for customers who travel frequently and use mobile devices to communicate with clients. To reduce the amount of false alerts, a change in location will be considered acceptable in such cases. However, this will make it more difficult to detect actual fraudulent activity.

Administrators can define a time interval during which end users are not supposed to make calls from different locations by configuring the **Location change allowed every** option in the product definition.

For example, a location change is allowed every 240 minutes (i.e. every 4 hours). A call comes in and its location is detected as "New York." For three days calls come from this number in New York, and then the account calls from Egypt. Merely an hour later, the same account calls from Bangladesh. Because only an hour had passed after the call from Egypt was made, the system considers the call from Bangladesh to be a potential security breach and screens it.

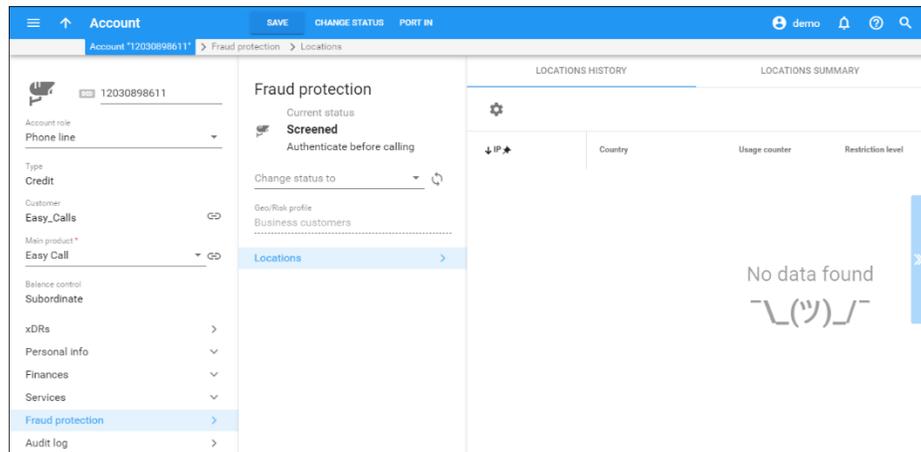
This restriction works independently of the geo / risk profile assigned, and applies to location changes between countries as well as within a single country. It applies to roaming and stationary end users (the latter are allowed to change location within the country specified in the **Current Location** option for the customer / customer site / account).

Two additional conditions may be introduced for an account:

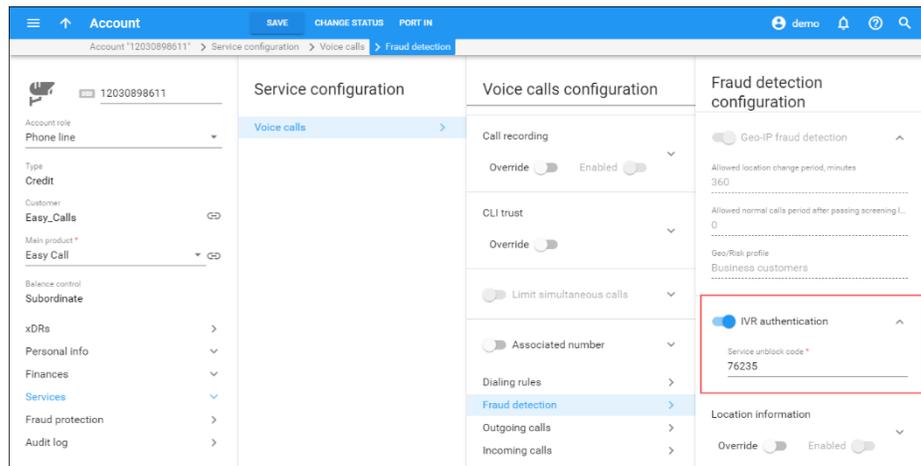
- **Screened** – This means that some unusual activity has been detected for this user. An attempt to make an outgoing call will connect the user to a screening IVR, where he or she must provide additional credentials to prove that this is indeed a legitimate user.

- **Quarantined** – This means that after being screened, this account was unable to supply valid credentials and is still generating a large number of call attempts. In order to reduce the load on your network, all such call attempts will automatically be blocked.

Every account is assigned a “default” location. An administrator may override this country selection in order to allow a legitimate user to keep using the service in a country that is otherwise considered “unsafe.”



Finally, each account is assigned its own unique “service unblock” code. This is provided to the customer upon sign-up, and can be used later to confirm that a legitimate customer is attempting to make a call, therefore allowing that call from a “suspicious” location.



Call filtering

For each outgoing call, the system analyzes the combination of the account’s status, its default location, and the location from which the call is made.

- Any call from a quarantined account is immediately rejected.

- A call from a screened account is redirected to the screening IVR.
- If a call is made by a stationary account but the country of its current location does not match the default country, the call will be screened. If a stationary account does not yet have a default country assigned (e.g. a newly created account), the call is allowed *only* if the current country location is not on the High-risk list.
- For roaming accounts, the call is allowed if the current country location matches the default country, or if the current country location is on the No Restrictions or Suspicious list.
- In all other cases, the call will be screened.

Screening IVR

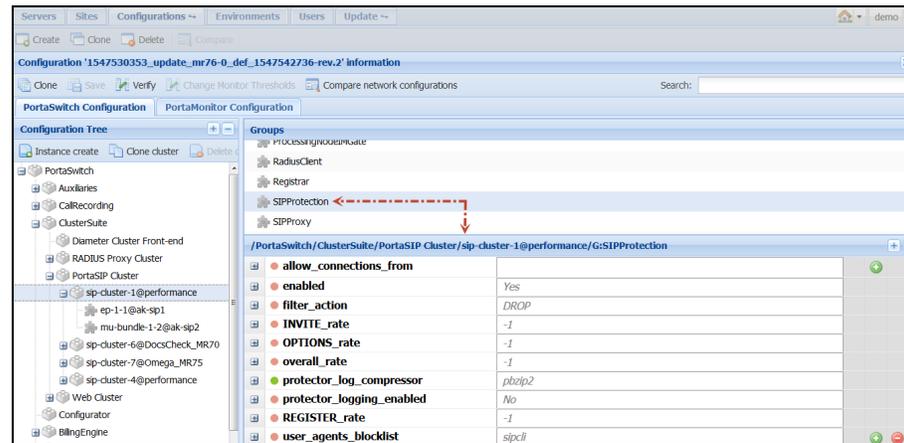
Clearly, blocking outgoing calls on any suspicion will create a negative experience for end users. A genuine customer may take his SIP phone on an overseas trip to see his family, or a businessman may need to make a very important call while transiting an airport in an “unusual” country. The screening IVR is designed to allow legitimate customers to continue making calls in situations where their activity is considered suspicious.

- The screening IVR announces that the user's call cannot be completed because of a potential risk of unauthorized usage.
- It asks the user to re-enter the random 3-digit code he will hear. This is done to confirm that there is a live person on the line. (Hackers often use auto-dialers to test vulnerabilities in a network or generate calls to premium numbers.)
- Then the user must enter the “service unblock” code provided to him when he signed up for the service. If the unblock code is entered correctly, the system will automatically connect the user to the originally dialed number (there is no need to re-enter the number).
- If there are multiple unsuccessful attempts in the screening IVR, the account will be switched to “quarantined” status. This (along with the 3-digit “live human” verification) prevents a hacker from using a script to find the unblock code by “brute force” and continue sending traffic.

Protection from DoS attacks

As Denial-of-Service (DoS) attacks are fairly common in the VoIP world, service providers must take proactive measures to ensure that service is not affected. The PortaSIP® DoS prevention feature utilizes the PortaSIP® cluster protector – the built-in firewall – to block network traffic coming from specific IP addresses once too many requests in a row have been sent (beyond the reasonable amount that would be generated by a legitimate SIP phone, proxy or gateway). The management interface

for the module is very easy to configure via the Configuration server web interface.



You may configure what maximum number of SIP requests will be accepted from each IP address, or set a limit for INVITE / OPTIONS / REGISTER SIP requests, separately. SIP requests above that limit will be dropped and logged for future investigation.

DoS prevention functionality also covers CPE provisioning via HTTP and TFTP protocols and is enabled for every web server instance by default. On the Configuration server, the administrator configures the time interval and the maximum number of packets that can be sent from a particular UA during that interval. If the UA sends too many requests for provisioning, any packets above that limit will be dropped.

Toll fraud prevention

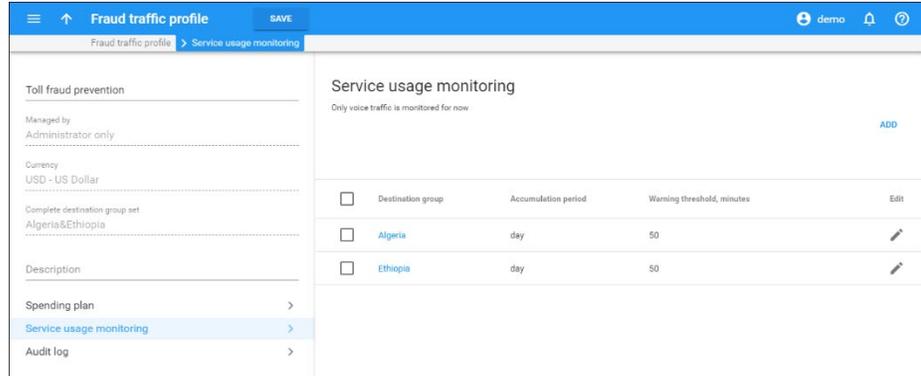
Toll fraud becomes an issue when a hacker breaks into a company phone system and routes unauthorized calls to international (long-distance) numbers. This type of fraud cannot be stopped by merely compiling an IP blacklist and, when it occurs, it usually results in substantial losses for the company. That is why it is important that administrators take timely measures to prevent it.

PortaBilling® offers a solution that promotes early stage toll fraud detection and sends alerts whenever any potential toll fraud activities take place.

An administrator defines the destinations that will be monitored for toll fraud (for example, Algeria and Ethiopia) by configuring a complete destination group set.

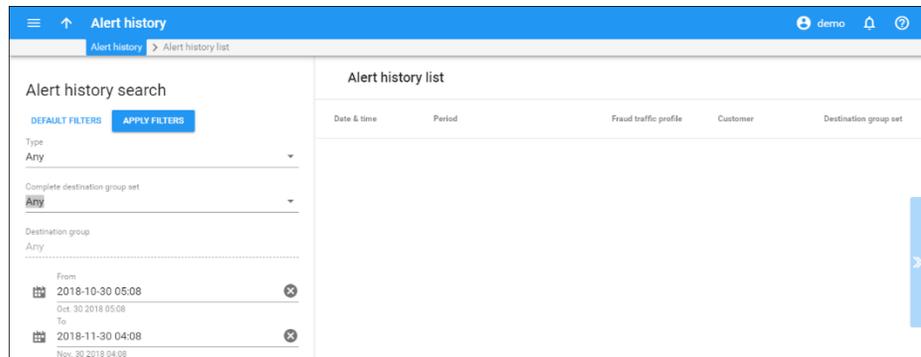
NOTE: Destinations from the destination group set must fully match the ones in the vendor tariff. Otherwise, the destination will not be monitored.

Then the administrator creates a **Fraud traffic profile** and selects the previously created complete destination group set for it. For each destination group in the set, the administrator configures a threshold that defines call durations during a span of time that are considered normal (for example, 50 minutes per day).



The administrator assigns the fraud traffic profile to a customer class or directly to a customer. If required, the administrator can override thresholds configured in the fraud traffic profile for an individual customer (for example, if sending larger / smaller volumes of traffic to monitored destinations is considered normal for this customer).

Whenever the duration of calls exceeds the set threshold, an alert is sent to the administrator. The administrator can review these automatically generated alerts on the **Alert History** page. To locate a specific alert, the administrator can filter the alerts by customer, customer class, destination group and time interval.



The same approach is used for early detection of unauthorized calls made to premium-rate numbers.

Fraud traffic profiles

Fraud traffic profile is an effective tool for minimizing customers' losses from fraudulent events and for monitoring customers' voice service usage.

Fraud traffic profile includes:

- Spending plan – This allows to limit customers' / resellers' daily expenses on voice calls and receive alerts once a specific limit is reached.
- Service usage monitoring – This allows to set the expected volume of voice traffic to be sent by your customers to unusual destinations (e.g. Somalia or Albania, while VoIP services are mostly used for the USA and Canada) and receive alerts once a threshold is reached.

An administrator configures a fraud traffic profile and assigns it to a customer / reseller via a customer class or directly.

Name	Managed by	Threshold	Currency	Check remaining funds	Complete destination group set
100 USD	Administrator only	100.00	USD		
50 USD	Administrator only	50.00	USD		Europe
Toll fraud prevention	Administrator only		USD		Premium-rate nu...

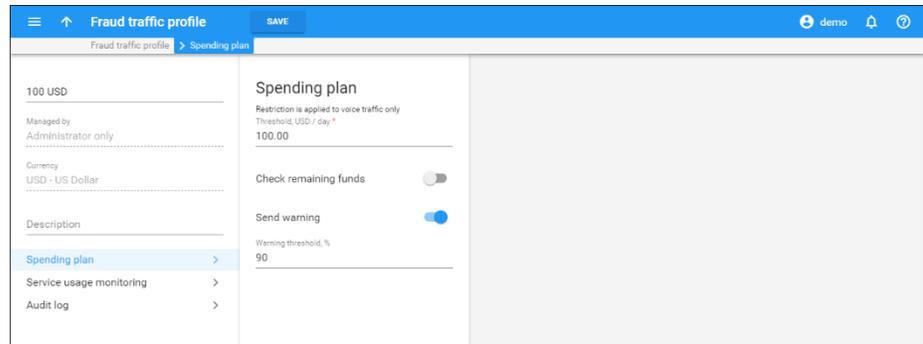
The administrator can also change a fraud traffic profile for a customer class. As a result, the fraud traffic profile will be changed for all of the customers / resellers with this customer class assigned.

Spending plan

Spending plans help service providers minimize losses from fraudulent events.

A spending plan defines an amount of money per day that a customer can spend on calls. PortaBilling® tracks a customer's daily charges, and increases the spending plan counter by the cost of every call completed.

Once a customer reaches their spending limit, their services are suspended. However, the customer can still make calls to emergency numbers. Regular service usage is reinstated on the following day, once the spending plan is renewed. The renewal time is midnight (00:00:00 in the customer's billing time zone).

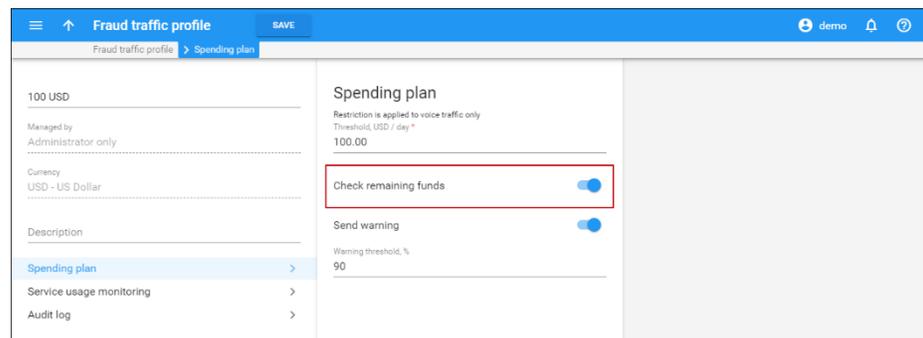


Note that PortaBilling® adds the full call cost to the spending plan counters *even if* a call is covered by a discount. To illustrate this, let's say a customer's quota includes 100 minutes of free calls to France. The regular price defined in the tariff is \$0.50 per minute. The customer makes a 5 minute call to France and is charged zero; however, their spending plan is increased by \$2.50 ($\$0.50 * 5$ minutes).

For each spending plan an administrator can define warning thresholds for receiving notifications when these thresholds' limits are reached. If a customer has used up their spending plan and wants to continue using the services, an administrator can temporarily increase their spending limit. End users can obtain more information about their spending plans on the customer self-care interface.

Spending plan remainder for call authorization

PortaBilling® can consider a customer's remaining spending plan amount during call authorization and re-authorization. If the remaining spending amount is lower than the available funds / credit limit, PortaBilling® uses it to authorize calls. This feature allows you to adjust the maximum permitted duration of calls in accordance to customers' daily limits, thus minimizing their potential losses in case of fraudulent activities.



The following example illustrates how it works:

Your SIP trunking customer EasyCall is assigned a \$100 spending plan. EasyCall has a credit limit of \$1000 USD and a balance of \$200 and they

are allowed to establish 10 simultaneous calls. When authorizing each new call, PortaBilling®:

- checks the customer's available credit amount (\$800) and compares it with the remaining spending amount (\$100).
- since \$100 is less than \$800, PortaBilling® calculates the call duration based on the remaining spending amount (\$100).

Let's say that a \$10 call has just ended. PortaBilling® produces a \$10 xDR and applies it to that customer. It also updates the remaining spending amount which now becomes \$90. When PortaBilling® authorizes a new call, it calculates the maximum call duration based on this remaining \$90.

When any of the calls in progress are re-authorized, PortaBilling® checks the remaining spending amount and calculates the call duration based on that.

Every time a call ends, the remaining spending amount decreases. Accordingly, PortaBilling® reduces the maximum call duration when it authorizes and re-authorizes calls. Once the spending limit is reached, PortaBilling® rejects all following authorization / re-authorization requests. Calls in progress last only for their previously authorized time.

Note that PortaBilling® does not lock funds from the remaining spending amount when it authorizes calls. Therefore, customers may still exceed their spending limit if a call in progress lasts too long and has therefore high cost by the moment the spending limit is reached. For example, if the customer has only \$1 left out of the \$100 spending limit and the call for \$10 ends, their total amount spent becomes \$109.

Spending plan and overdraft protection settings

The remainder of the spending amount can override the overdraft protection configuration if it is less than the amount to be locked for a call.

Let's say that an administrator has configured PortaBilling® to lock at least \$10 per call and the customer is allowed to call A-Z destinations for a unified price of \$1 /min. The customer's remaining spending amount is now \$8.

When a customer makes a new call, PortaBilling® compares the amount to lock (\$10) with the remaining spending amount (\$8). Since \$8 is less than \$10, PortaBilling® reduces the amount to be allocated for the call by up to \$8. PortaBilling® then calculates the maximum call duration based on \$8. Thus, it authorizes the call to last eight minutes and therefore locks \$8.

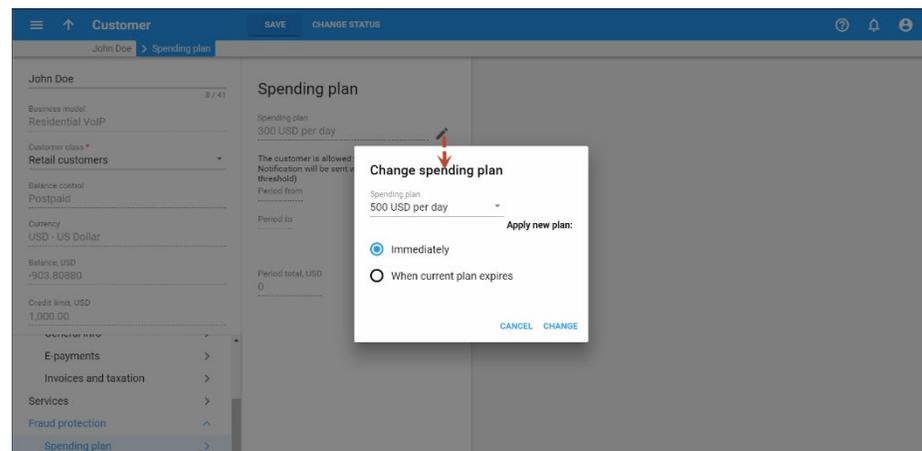
Deferred spending plan change

When an administrator changes a spending plan for a customer (for example, from spending plan A to spending plan B), within the new spending plan they can include the amount already spent.

For example, if customer John Doe has a spending plan of 300 USD and has already spent 250 USD, an administrator can change John Doe's spending plan to 500 USD. Now John Doe can spend 250 USD more (500 USD - 250 USD = 250 USD).

To do this, the administrator needs to choose a desired spending plan and then select the **Apply the new spending plan immediately** option. This gives the following results:

- The spending limit (the amount of money the customer may spend on services per day) is changed according to the new spending plan.
- The amount of money the customer has already spent is not reset and continues being calculated.
- The spending plan expiration date and time are not changed. So a customer can use the new spending limit until the previous spending plan expires.



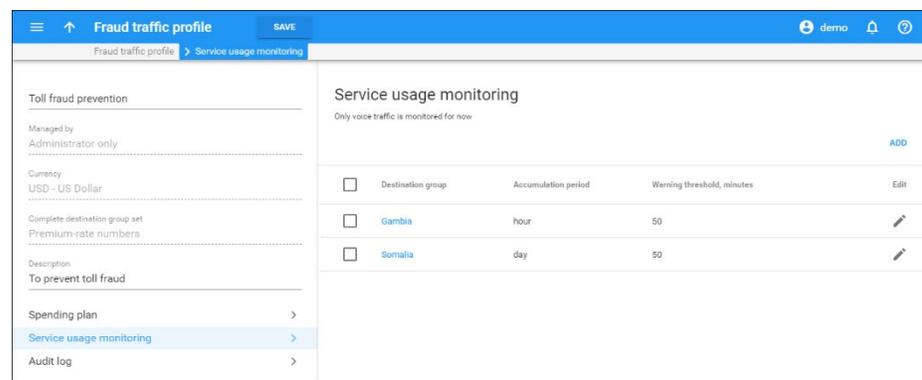
If the administrator selects the **When current plan expires** option, the new spending plan will take effect at midnight (00:00:00) after the previous one expires. The customer's spent amount will be reset to zero and calculated from scratch towards the new spending limit.

This gives administrators more flexibility for adjusting customer's spending plans. Having the capability to change a spending plan for a customer class allows administrators to apply these changes to many customers at once, thereby reducing the administrators' workload.

Service usage monitoring

Service usage monitoring allows service providers to monitor traffic sent through their networks and receive alerts whenever a toll fraud calling pattern is detected.

For example, for the service provider who provides VoIP services mostly for the USA and Canada, an administrator can configure the PortaBilling to monitor international traffic and send alerts when an unusual burst of international traffic is detected. For the service provider who normally has lots of international traffic, the administrator can configure the service usage monitoring to receive alerts when an unusual amount of traffic is sent to premium-rate numbers.



This allows service providers to promptly detect suspicious calling activities and take the necessary measures to mitigate potential losses.

Fraud traffic profiles for resellers use

Fraud traffic profiles can only be used by resellers to track their customers' daily expenses on voice calls and notify a reseller when a customer's spending threshold is reached. Therefore, resellers see their fraud traffic profiles as spending plans on their self-care portal.

A reseller creates a fraud traffic profile (i.e. a spending plan) and assigns it to their customer via a customer class or directly. PortaBilling® then tracks the customer's daily charges and once the spending limit is reached, this customer's services are suspended.

If a customer has used up their spending plan and wants to continue using the services, the reseller can temporarily increase their spending limit. End users can obtain more information about their spending plans on the customer self-care interface.

For each spending plan the reseller can define warning thresholds for receiving notifications once a customer crosses their spending threshold.

To enable resellers to receive notifications, the administrator enables notifications for a customer class. Then, all resellers / subresellers who are assigned this customer class are notified.

NOTE: An administrator does not receive notifications about reseller's customers, and a reseller does not receive notifications about their subreseller's customers.

Binding a user session to an IP address

Session hijacking is a hacker practice: a fraudster obtains a valid session ID and uses it to gain unauthorized access to the system. The SSL encryption used by PortaSwitch® provides adequate protection against most of these kinds of traffic interception attacks. However, there might still be a slight possibility that a hacker can guess the session ID or that, while travelling through a high-fraud-risk country, an end user may connect to an unscrupulous internet provider that has the means to decrypt traffic and leaks sensitive information to fraudsters.

Binding a user session to an IP address protects PortaBilling® web interface users from the above-described risks.

This is how it works:

When a user logs in to the web interface, the created session is associated with the user's current IP address. The system subsequently discards all access attempts from other IP addresses within that session. It also sends a notification to the user when such an attempt occurs.

Thus, even if someone manages to steal a session ID, it is of no use to the fraudster.

In the case of a legitimate IP address change (for example, the user has reconnected to the Internet via a different WLAN), the user sees a notification page with information on how they can continue their work. The user has two options: either authorize their new IP address by following the link provided in the notification email, or simply re-login.

To enable the binding of a user session to an IP address, an administrator must go to the Configuration server web interface and set the **Security.EnableSessionIdProtection** option to **Yes**.

It is possible to additionally configure the following options in the **Security** group on the Configuration server web interface:

- How many IP addresses can be authorized within the same session.
- After how many access attempts from an unknown IP the system will reject subsequent ones without any notification.

10. Internet provisioning and charging

Broadband Internet access service

PortaBilling® is a versatile tool, and the range of services that can be billed extends far beyond the realm of VOIP. With many VOIP providers also (or primarily) being ISPs, billing Internet access is naturally high on the priority list.

Supported services

PortaBilling® supports various billing models for Internet access:

- Fixed bandwidth with unlimited amount of traffic, monthly fee.
- Fixed traffic (quota) at maximum bandwidth, monthly fee.
- Per-megabyte charges at fixed or maximum bandwidth.
- Per-minute charges at fixed or maximum bandwidth.

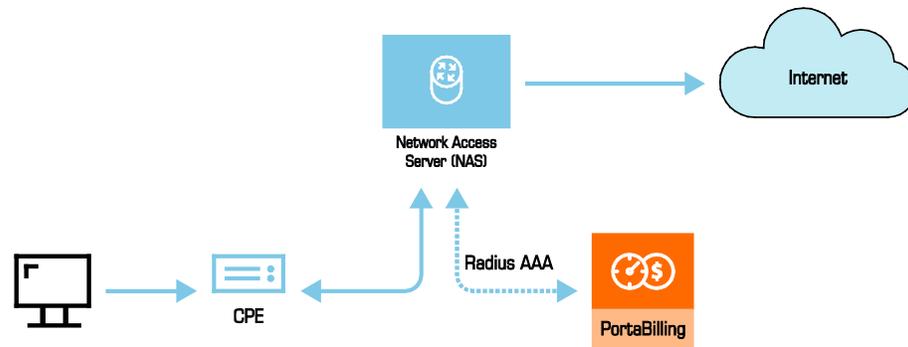
All the models are supported when PortaBilling® communicates directly with a Network Access Server (NAS). In this case, it acts as an AAA, including bandwidth/quota control, accumulating usage statistics for user reference, and disconnecting users upon suspension, reached quota or account expiry.

There is also an option to use third-party service controllers (bandwidth/quota control), which are automatically provisioned from PortaBilling®. Only first two models are supported in this case.

Provisioning

PortaBilling® allows you to manage the provisioning of Internet access services when a user establishes a session with your NAS (Network Access Server) for further Internet access. PortaBilling® performs the session authorization (validating that this user is allowed to use this service and what the maximum allowed session time is) and, when the connection is terminated, deducts the cost from the customer's balance according to the session time and the specified rate (for the special destination **NETACCESS**).

NOTE: As a good practice we recommend that you configure your equipment to send keep-alive messages to PortaBilling® at least every 2 hours while you are using Internet access services. Note that PortaBilling® applies a single rate for a session between two keep-alive messages, even if it spans peak and off-peak periods. Therefore, to ensure billing accuracy, make keep-alive intervals shorter for scenarios where there are different rates involved.



1. A CPE device (for instance, a DSL modem or PPPoE client on a desktop computer) connects to a NAS (gateway) requesting that a new session be established. The protocol used for the session between the CPE and NAS (e.g. PPP or PPPoE) may differ depending on the access technology used, but from the billing perspective this is irrelevant.
2. The NAS forwards account information to PortaBilling® for authentication using RADIUS. Porta Billing verifies the account information:
 - credentials (“is the account valid?”),
 - balance/quota (“does the account have enough funds/traffic quota/minutes?”),

and returns the result to the NAS.

Additional information regarding the session (e.g. static IP address) may be returned from the billing to NAS.

3. If the account is successfully authenticated, the NAS establishes the session so the end user can access the Internet.
4. While the connection is active, the NAS sends periodic updates about the amount of data transferred within the session, and PortaBilling® repeatedly checks the validity and balance of the account (re-authorizes it). If the account has a megabytes/minutes quota, PortaBilling® instructs the NAS to disconnect the session as soon as the quota is reached.

PortaBilling® can be used as a back-end for any NAS which uses a Cisco-compatible RADIUS protocol.

Controlling bandwidth

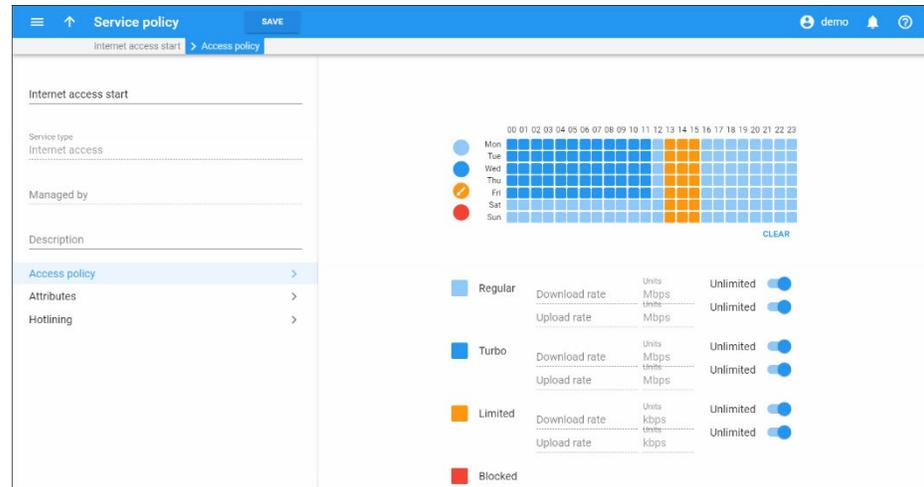
PortaBilling® allows you to control the amount of Internet bandwidth allocated to various types of customers or to even override the default bandwidth policy for an individual customer.

In order to better oversee network utilization, it is very important for an ISP to control the bandwidth available for specific types of customers throughout the day. This includes, for instance, limiting the amount of

bandwidth consumed by residential customers during business hours, thus ensuring that business customers receive adequate service; or creating special products for accessing the Internet at maximum speed during nighttime hours, when the network is underutilized.

Internet access policy

A policy is a collection of rules and parameters which define how an end user is allowed to access the Internet.



There are four available types of Internet access:

- Regular access (default).
- Turbo access – Access during off-peak hours.
- Blocked service (the customer is not able to access the Internet at all).
- Limited access. Typically this is used as an alternative to “blocked” in situations where a customer does not have sufficient funds or failed to pay his last invoice on time. While the customer will not be able to surf the web or download normally, he can still send or receive emails and use the customer self-care interface to submit payment.

Each policy includes:

- A scheduling table, which designates what type of access is used during each time period.
- Specific values for the allowed upload / download speed for each type.

When a new period starts, PortaBilling® sends a request to NAS to change the parameters of the current session. This is done using CoA (Change-of-Authorization), for equipment which supports it (in this case the bandwidth re-allocation is done transparently to the end user), or POD (Packet of Disconnect, which terminates the current session).

Rating data

Since Internet access is an “always-on” service, billing is typically done in real time, based on RADIUS update (“keep-alive”) requests received from the NAS; thus many keep-alive requests may be combined into a single xDR. In order to provide detailed information on the charges applied, the contents of each update request are stored in the PortaBilling® internal tables, and both summaries and detailed data can be browsed on the PortaBilling® web interface.

Data transfer quotas

Data transfer quotas function in a similar way to volume discounts for services such as voice calls. You can define combinations of data transfer thresholds, along with the applicable discounts, e.g. the first 5GB free (100% discount), after which the customer is charged the normal rate (\$0.01/MB).

You also have the ability to adjust the service itself once a threshold is reached, e.g. when the customer exceeds 5GB his download speed will be severely reduced (but he can still use the web, email and the like). If he continues to use the service in an abusive manner, and so exceeds 6GB, the service will be automatically blocked.

Service shutdown

PortaBilling® can initiate the disconnection of a session in progress on NAS; for instance, if it is determined that the customer has exceeded the allowed data transfer limit, or if the customer has been suspended due to non-payment. In this case, PortaBilling® sends NAS a POD (Packet of Disconnect).

Rating configuration

WiFi

The rate for this special destination will be used to charge your customers for WiFi access service.

NETACCESS

The rate for this special destination will be used to charge your customers for Internet access service such as DSL, PPPoE or WiMAX.

Routed network provisioning

In cases when a customer needs to purchase an entire IP-network from the ISP, PortaBilling® provides the option for configuring the routing information for the customer's IP-networks on the NAS. To do this, the administrator must indicate the network prefix, the length of the network prefix mask and the gateway IP address in the **Routed Networks** attribute. This routing information will then be automatically propagated to NAS each time the customer connects.

Internet service provisioning for ADSL users via Cisco 7000 / 7200 series routers

With Cisco 7000 / 7200 series routers ISPs can change Internet speed without interrupting users' active sessions, depending on the time-period, and redirect users to a hotline portal for ADSL users.

Seamless Internet speed change

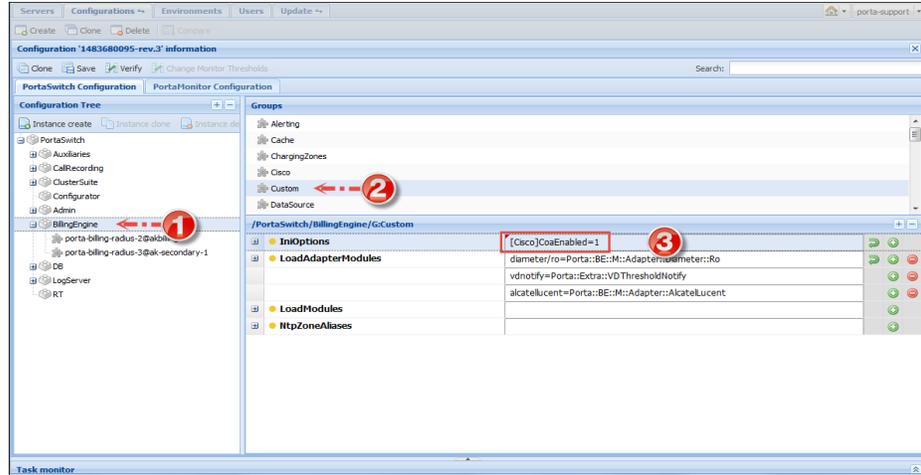
When a user establishes an Internet session, the NAS sends an authentication request to PortaBilling®. Upon session authentication, PortaBilling® generates a response defining the time-period for the session start.

When PortaBilling® detects that the time period has changed, it sends a CoA (change of authorization) request to the NAS. The NAS then adjusts the Internet speed for the current session, according to the Internet access policy parameters, without terminating the session.

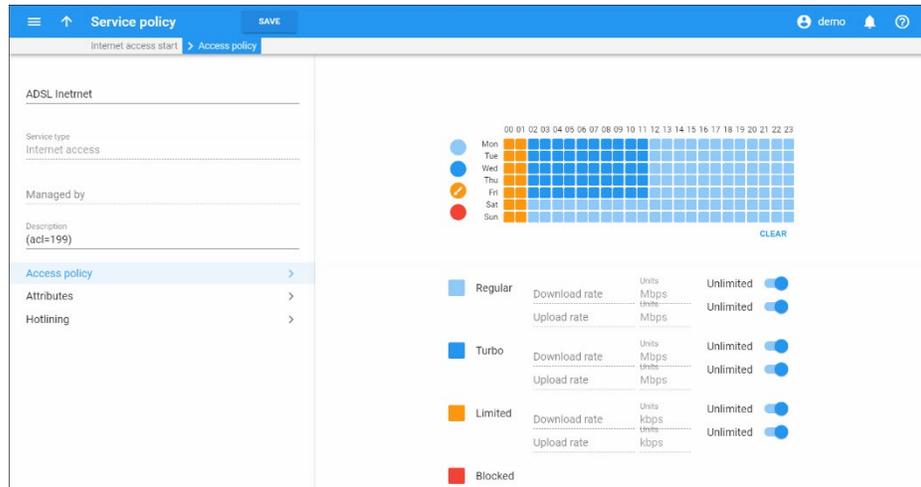
Configuration

To configure this functionality, follow the next steps:

1. On the Configuration server add `[Cisco]CoaEnabled=1` value for the **IniOptions** option.



2. On the administrator web interface, define time-periods and download / upload speed on the **General Info** tab for the Internet access policy.
3. Specify ACL in the description field (“(acl=199)” is a recommended value for the ACL).



4. Assign the policy to the account.

This provides ISPs with more flexible Internet provisioning and allows the Internet speed to be changed without session interruption.

Hotlining

Cisco 7000 / 7200 Series routers allow ADSL users to be redirected to a defined web site when:

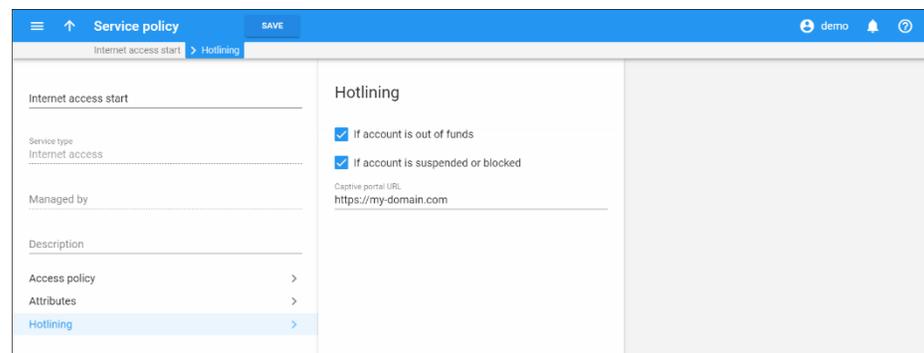
- their Internet access is blocked according to the time-period set in the Internet access policy, or
- their resources have all been consumed.

During the active Internet session PortaBilling® detects that the time period has been changed to “blocked.” It sends a PoD (Packet of Disconnect) request to NAS. The NAS then disconnects the current session and sends a re-authentication request to PortaBilling®. PortaBilling® detects that the service is blocked according to the Internet access policy and generates a response. NAS then redirects the end user to a defined website.

Sometimes user’s funds become depleted or their quota is reached while still using the service. In these cases, once the lack of resources has been detected, PortaBilling® sends the NAS a PoD (Packet of Disconnect) request. Cisco / NAS then disconnects the session and attempts to re-authorize the user. Having received a re-authorization request, PortaBilling® detects the lack of funds for the service and instructs the NAS to redirect the user to the hotline portal. The NAS then redirects the end user to a defined website.

Configuration

An administrator enables the feature on the **Hotline** page of the Internet access policy and assigns it to the account.



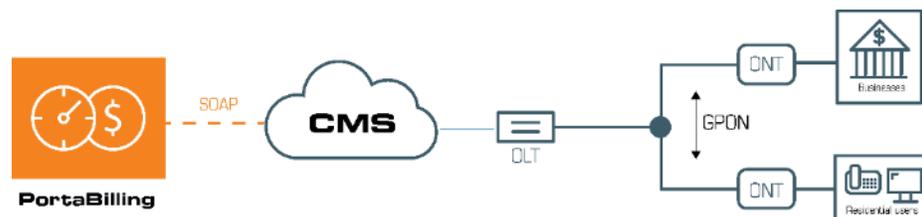
With these features, ISPs introduce advanced service provisioning to their users.

Fiber to the home service provisioning via Calix ONT

Gigabit passive optical network (GPON) and Active Ethernet (AE) are fiber access technologies that provide gigabit Internet services for subscribers. You can use them with PortaBilling® and the Calix network access equipment to deploy fiber to the home (FTTH) networks and deliver high-speed broadband Internet to your subscribers.

GPON / AE networks consist of:

- OLT – Optical Line Terminal – the service provider endpoint, and
- ONT – Optical Network Terminal – the endpoint that is located on the customer premises.



PortaBilling® provisions Calix ONTs such as:

- Calix 803G GigaPoint;
- Calix 844E GigaCenter;
- Calix 844G / 854G GigaCenters;
- Calix 716GE-I.

and Calix E7-20 EXA Service Access Platform for GPON / AE services.

To make this happen, PortaBilling® has integrated with the Calix Management System (CMS). PortaBilling® provisions the following information to the CMS via the External Systems Provisioning Framework (ESPF):

- ONT data such as the FSAN serial number – e.g. the unique serial number of the equipment, and
- account and service configuration data, e.g. the account status, Calix bandwidth profile ID within the Internet access policy, etc.

Thus, Calix performs network access and configuration management while PortaBilling® operates as a B/OSS by providing a single place to perform customer and service management.

When the administrator assigns the ONT to a user's account from the CPE inventory in PortaBilling®, PortaBilling® sends the provisioning request to the CMS. The CMS activates the service for the user on the

ONT and allocates the bandwidth as defined in the bandwidth profile. The user can now access the service.

When a user does not pay for their service (e.g. has insufficient funds to cover the monthly subscription fees), PortaBilling® instructs the CMS to suspend their services on the ONT until a payment is received. If the user upgrades to a new package (e.g. from 5Mbps to 20Mbps), PortaBilling® provisions the service configuration change to the CMS. The CMS allocates the new bandwidth limits on the ONT and the user receives the benefit of faster speed.

To configure the Calix ONT provisioning, do the following:

1. Configure the Calix Management System by creating network uplinks for GPON services, ONT and Ethernet bandwidth profiles, etc. (Please find configuration details on the [Calix](#) website.).
2. Upload ONT devices to the CPE inventory;
3. Configure the Internet Access policy and specify the Ethernet bandwidth profile ID within.
4. Configure the Internet access service in PortaBilling®.
5. Configure the ESPF in PortaBilling to enable data provisioning to CMS.

This solution enables you to build cost-effective last mile networks that work with gigabit speeds and thus meet users' increasing demands for more bandwidth. You automate the service configuration and the delivery to your users as well as its availability, based on their demands (e.g. the decision to upgrade), and their billing status (e.g. suspended upon non-payment).

Provisioning of Calix ONTs in residential gateway mode

You can provision a Calix ONT (Optical Network Terminal) in residential gateway mode (RG). Residential gateway mode enables you to limit bandwidth on the entire ONT, rather than on a specific Ethernet port. As a result, you set speed limits per household.

You can create different packages such as 500Mbps, 1Gbps, etc., to meet your customers' needs. For example, you can offer higher speeds to premium customers and enterprises.

PortaBilling® is integrated with the Calix Management System (CMS) for ONT provisioning. When an administrator creates a new user, PortaBilling® provisions the bandwidth profile ID, the account status, etc. In addition to account and service configuration data, the ONT WAN port type is provisioned to the CMS. This port type is defined as RG.

To provision a port type to Calix CMS, an administrator configures the following in PortaBilling®:

- creates the Internet access policy that selects RG for the ONT slot type;
- assigns the policy to the product;
- assigns the product to a user's account.

The screenshot displays the 'Service policy' configuration page for a GPON service. The breadcrumb trail is 'GPON > Attributes > Calix'. The left sidebar shows the 'GPON' service type with various configuration options like 'Service type', 'Managed by', 'Description', 'Access policy', 'Attributes', 'Hotlining', and 'Audit log'. The main content area is divided into three sections: 'Attributes', 'Calix', and 'Calix 2 enabled attributes'. The 'Attributes' section lists vendors such as Alcatel-Lucent, Cisco ASR, Cisco NAS, Cisco PDSN, ECI, Gy, Huawei, Huawei BRAS, and Juniper. The 'Calix' section is expanded, showing 'Bandwidth profile id' set to 15 and 'Data ont slot type' set to RG. The 'Data ont slot type' dropdown menu is highlighted with a red box.

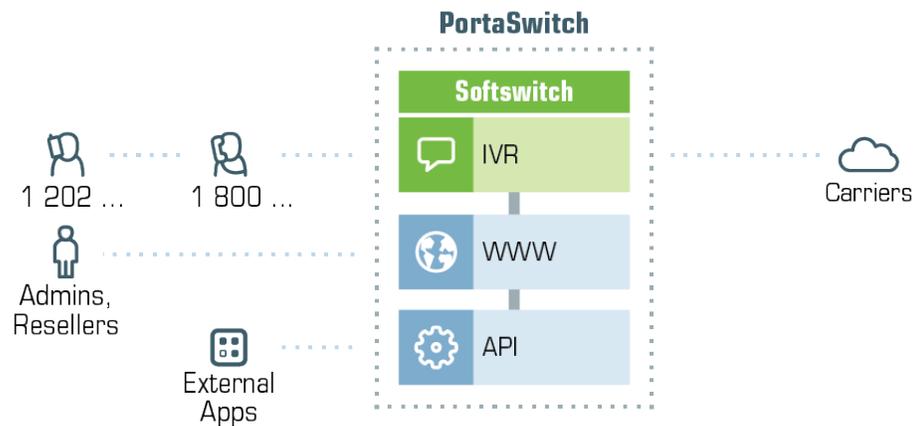
Note that only GPON ONTs such as Calix 844E GigaCenter, 803G GigaPoint, etc. can be defined as residential gateways.

The WAN ports are provisioned as Gigabit Ethernet port types, by default. You can change the RG port type in the service policy or CPE profile. The configuration in the CPE profile takes precedence over the service policy configuration.

Provisioning the ONTs in residential gateway mode helps you deliver Internet services to residential users and effectively manage bandwidth limits for your customers. As a result, you remain competitive in the market and prevent revenue leakage.

WiMAX services

In the case of WiMAX, PortaBilling does not directly communicate with any of the base stations; instead it communicates via the RADIUS protocol with WiMAX ASN-GW.



User authentication

PortaBilling supports WiMAX user authentication via EAP-TLS (this method relies on the certificate provided in the CPE) or EAP-TTLS (this method requires the CPE to have a valid username / password combination). Authentication also includes generation of the WiMAX session key and other required processing.

Flow provisioning

PortaBilling may return optional attributes to provision service flows for the user's WiMAX session.

Bandwidth provisioning

The administrator may define the allowed download/upload bandwidth as part of the product configuration in PortaBilling. Then during the service flow provisioning mentioned above the required information will be returned to the ASN gateway, and so will be applied to the session.

Hotlining

If a user is not allowed to use the WiMAX service (insufficient balance, customer blocked, account's product does not include the WiMAX service, or some other problem), instead of returning an authentication reject PortaBilling may allow the session to be established, but will assign a specific hotline profile to it. When the end user tries to browse the Internet, he will automatically be redirected to a separate web page informing him about the status of his account.

Session monitoring

When a WiMAX session becomes active, ASN-GW sends a start accounting request to PortaBilling, informing it that the customer is now online. This information is stored in PortaBilling, so that the administrator

can see information about currently active sessions on the PortaBilling web interface.

Real-time processing of service usage

When a customer is connected to a WiMAX network, ASN-GW periodically sends interim (also called *keep-alive*) accounting requests with information about the currently consumed amount of service – for instance, the total number of bytes downloaded. Upon receiving an interim accounting request, PortaBilling updates the internal tables for service usage (so the administrator and the end user can see up-to-date statistics) and re-calculates the applicable charges for this session. Thus the charges are immediately deducted from the available balance or (in case of quota enforcement) the session is disconnected almost immediately after the quota is reached.

Postpaid service

The postpaid model allows you to authenticate only those customers who are entitled to use the service and have their account in good standing. When a customer is blocked by the administrator, suspended because of non-payment, or exceeds his credit limit, a CoA (change of authorization) message will be sent to ASN-GW to terminate the current session or switch it to the hotlined state.

Prepaid service

Each new session is authorized based on whether the account has funds available for using the service. After receiving each interim accounting request, the charged amount for this session is recalculated and applied to the account's balance. If there are no remaining funds, PortaBilling sends a CoA request to ASN-GW to terminate the current session or switch it to the hotlined state.

Data transfer quota enforcement

The PortaBilling administrator defines data transfer policies using a flexible tool – volume discount plans. This allows you to define the amount of data provided free of charge (at a 100% discount) and what happens after that: the service can be blocked, switched to a “limited” mode (low bandwidth), or the customer can be charged for data above the threshold.

Supported equipment

PortaBilling supports the generic set of WiMAX RADIUS attributes and may potentially be used for basic authentication and accounting of WiMAX sessions with all major WiMAX equipment vendors. Provisioning of the flow attributes (or other proprietary attributes) requires vendor-specific interop. Please contact PortaOne to obtain more

information about the currently supported vendors and equipment models.

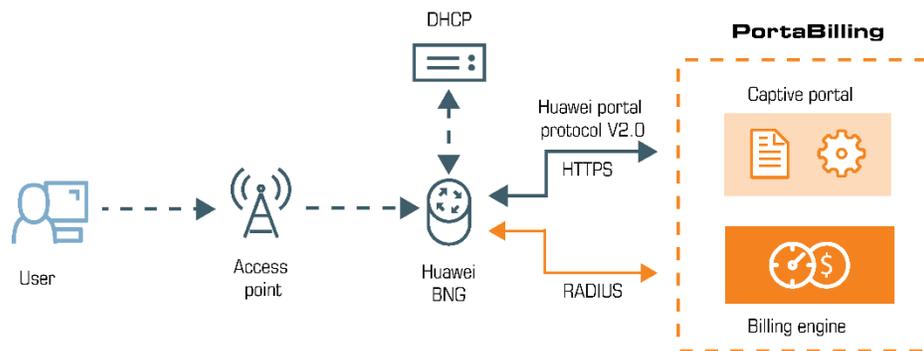
Captive portal for WiFi hotspots

PortaBilling® has a captive portal for WiFi hotspot services. This allows you to configure WiFi access points in various public places (e.g. cafes, gyms, hotels, etc.), enable users to easily subscribe and access WiFi using vouchers. To promote your service you can also configure PortaBilling® to grant one-time free minutes (e.g. 15 minutes of free WiFi access per day).

The supplied captive portal interoperates only with Huawei BNG (Broadband Network Gateway) and is a sample model. Your in-house or outsource developers can extend it to interoperate with the WiFi routers and access points that you use. Alternatively, they can use it as a basis when building your customized captive portal.

The captive portal is located on the PortaBilling® web server and consists of:

- A front-end web application – This is the set of webpages to which users are redirected when they open a web browser. These pages contain your legal information (name, logo and the terms of use), provide signup / login options and show available prepaid plans. They also return user credentials for subsequent logins upon signup, display the session status and the remaining time available. You can customize the content, page design and layout (e.g. use your pictures and colors) as you like.
- The back-end signup module. This communicates with PortaBilling® via the API to create user accounts. It also communicates with Huawei BNG via Huawei portal protocol V2.0 to pass the user credentials and their device IP address to PortaBilling® for authorization. Upon your request, other communication protocols can be added to integrate the signup module with another BNG or router you use.



To illustrate how the captive portal works, consider the following example:

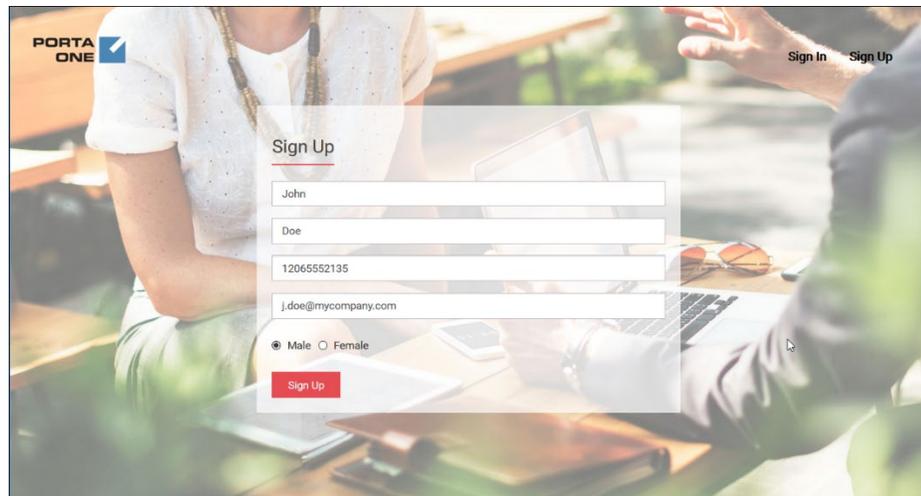
Let's say you provide the following WiFi access options:

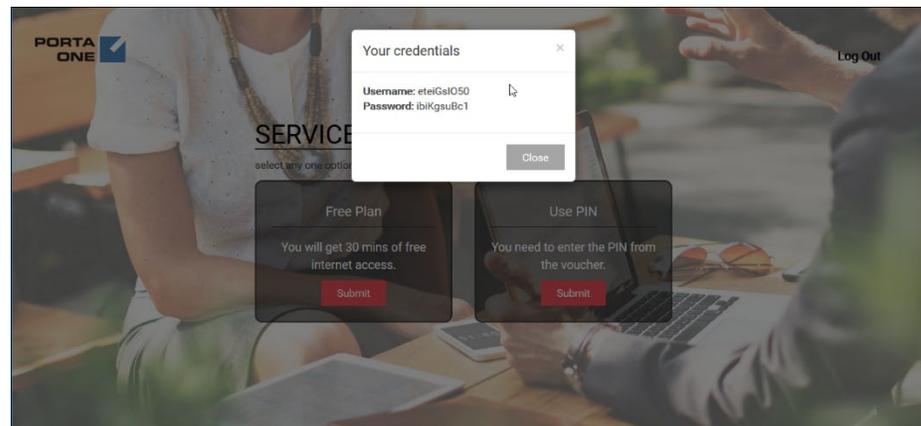
- 30 free minutes per day;
- \$10 vouchers for 1 hour; and
- \$15 vouchers for 2 hours.

John Doe wants to use the Internet, so he purchases a \$10 voucher. When he connects to WiFi, his device is assigned an IP address by the DHCP server. He opens the web browser and the BNG redirects him to the captive portal to accept the terms of use.

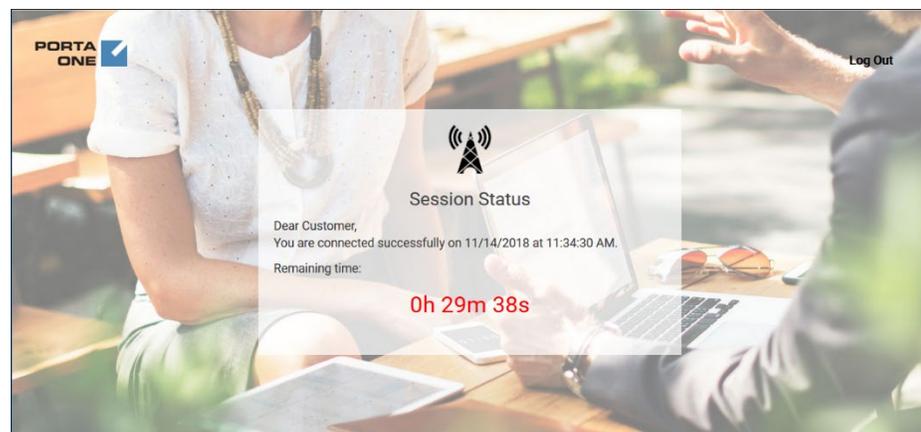


Then a sign up screen appears so he enters his name, email, gender and phone number. The signup module creates an account for him in PortaBilling® and the portal's webpage shows his credentials for further logins.





John can choose to use free minutes or enter the voucher PIN. He chooses the free minutes first. The signup module sends John's credentials and the device's IP address to the BNG. The BNG authorizes the account in PortaBilling®, receives the session duration time from PortaBilling® in the response and activates his Internet access.



When the free 30 minutes expire, the BNG sends the accounting information to PortaBilling® and redirects John to the captive portal. John must enter his voucher PIN now to use the Internet. When he enters it, PortaBilling® authorizes his account for one-hour access and the BNG resumes his session.

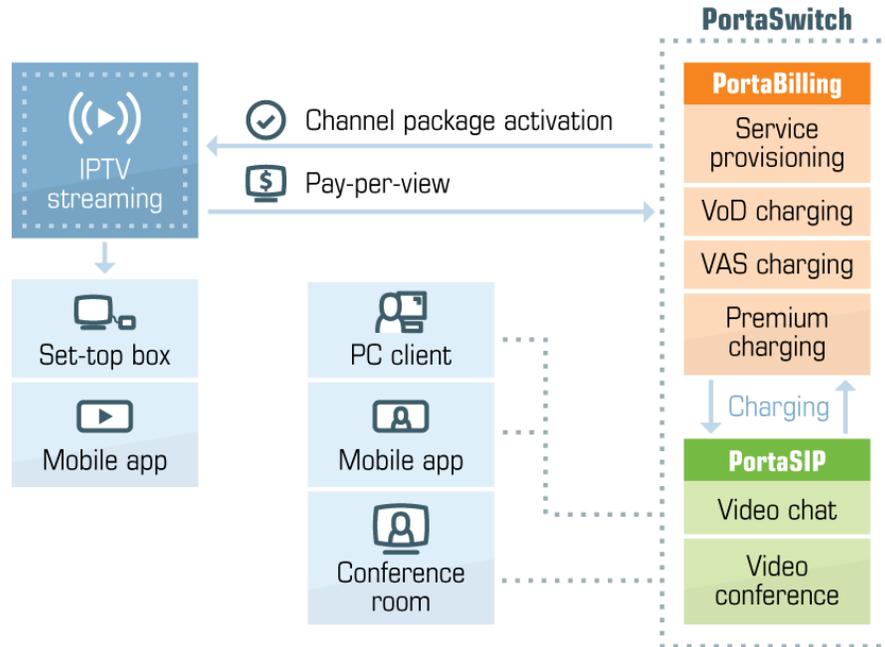
After 45 minutes, John closes the session. He still has 15 minutes of Internet usage available that he can use by logging in with the credentials provided.

Thus, with PortaBilling® and the captive portal you can generate revenue by selling WiFi vouchers and providing WiFi connectivity in public places like cafes, hotels, libraries, airports, etc. For how to configure the captive portal, please refer to the [PortaSwitch Handbook collection](#).

11. IPTV provisioning

IPTV

For today’s service providers, “triple play” is not just a buzzword - rather, it is the only way to stay competitive and retain their customer base. Combined solution provided by PortaOne in conjunction with an IPTV platform can be used to add IPTV to your service portfolio and expand your business growth.



An IPTV platform normally includes a content delivery network, a services framework, application servers and client portals. Together they enable the acquisition, storage, distribution, caching and delivery of content to a variety of devices; the management of content workflow; rights management and royalty reporting to rights owners; and the delivery of IPTV content to various applications (such as set-top boxes).

PortaBilling is a converged real-time billing system, service provisioning environment and CRM used for rating and invoicing any type of service you provide to customers, with centralized funds management and consolidated bills.

The administrator can define multiple channel packages for IPTV in the IPTV platform. These are then imported into PortaBilling® via XML API and a channel package can be assigned to the customer product together with a subscription plan for it (just like for any other service, such as internet telephony, messaging, or data transfer). This allows you to do

better marketing, work with multiple customer groups, and enjoy greater flexibility in terms of pricing the service.

When a new or existing customer signs up for the IPTV service, the available channel information (based on the assigned channel package) is provisioned into the IPTV platform via an extensible XML API. Also, if a customer's product configuration changes, or a customer is blocked or suspended (e.g. due to non-payment), this information is transferred to the IPTV platform.

Support for the pay-per-view service will be introduced in a future release, at which time when a customer requests pay-per-view content, it will be validated in real time – if he has enough funds available – and the cost is debited from his account once he confirms the purchase.

As a result, the IPTV service is fully managed via the PortaBilling® web interface, and all applicable charges are added to the customer's invoice, along with other charges (periodic and event-based), for services such as internet access, voice calls, and so on. This enhances the customer's experience, improves accounts receivable tracking, simplifies maintenance, and reduces operating costs.

The following IPTV platforms are currently supported:

- Matrix,
- Minerva,
- Telebreeze.

The combination of an IPTV platform and PortaBilling creates a world-class video ingest, storage, distribution and delivery platform with a robust, high-quality video functionality. Easy provisioning and full control on the billing side allow you to reduce operating costs and quickly achieve profits from IPTV services.

12. xDR mediation and import

xDR Mediation concept

Administrators can import CDRs to PortaBilling® from various external sources by using the xDR Mediation utility. The utility extracts the CDRs from incoming files and arranges the ready for processing CDRs into collections. PortaBilling® calculates the charges for these CDRs according to defined rates, and then based on this information, creates transaction records (also called xDRs). The created xDRs are then associated with corresponding accounts and vendors.

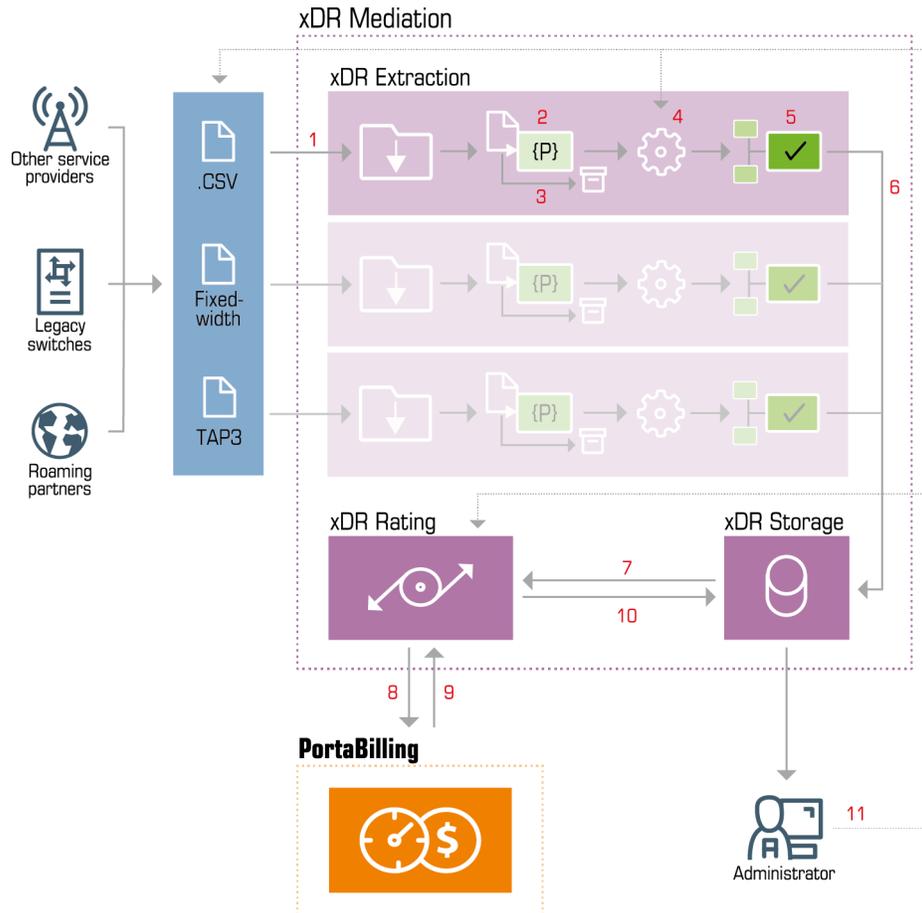
Service providers face the necessity to import CDRs into PortaBilling® when their telecom partners deliver CDRs to them in one of several file formats or when the legacy equipment used does not support RADIUS and can only save CDRs to files. The ability to import CDRs in such cases helps to ensure accounting and reporting integrity. It eases the limitations on the equipment that can be used and allows the service provider to extend the pool of potential business partners.

xDR Mediation capabilities

- CDRs can be imported to PortaBilling® from .csv, fixed-width, TAP3 files, and from ASN1 files generated by Huawei MSOFTX3000.
- CDRs can be imported in chronological order based on a time stamp. This is the time when a billing event takes place (e.g. the connect time for a voice call).
- It is possible to import CDRs for any types of services, for instance, voice calls, mobile and broadband data transfer or messaging.
- PortaBilling® calculates the charges for imported CDRs.
- Service providers may import xDRs that already contain charged amounts to PortaBilling®. In this case, import can be configured to forward these amounts to PortaBilling® as they are. Then PortaBilling® can apply an extra charge to the corresponding XDRs, if required.
- Using Perl code, an administrator can define tailored data manipulation rules to apply to each type of source data. Created rules may, for example, translate CLD into E.164 format. Complex transformations can be executed as a custom module.
- Several partial CDRs can be merged into a single logical CDR or one CDR can be split into several (i.e. to create separate call records for calling and called parties when calls are made within the same IP Centrex).

- An administrator can check in real time whether CDRs are imported and rated successfully. In case there are issues with any of them, the administrator can adjust the extraction or rating configuration and re-run the process.

xDR mediation process



Let's have a look at the xDR import process step by step:

- The input file with CDRs is saved to the incoming folder (1). When several files are added to the folder they are automatically placed in an alphanumerical order.
- The queue manager that constantly monitors the incoming folder sees the file, checks that it is not a duplicate and queues it for processing.
- Once queued, the Extraction utility creates an empty CDR collection – a set of CDR records from one or several files – in the interim data storage (referred to as collections' storage).
- The Extraction utility reads the file content and parses the data records (2). When parsed, each CDR is represented as a set of attribute-value pairs.

- The Extraction utility moves the processed file to the archive location and (optionally) compresses it (3).
- Then the Extraction utility performs the required data manipulation on the parsed CDRs according to the rules defined by the administrator (4).

```
User-Name=18005551155
Acct-Session-Id=01010101
Calling-Station-Id=18005551155
Called-Station-Id=18005555555
Acct-Session-Time=65
h323-connect-time=02:15:03.000 UTC Mon Dec 23 2014
h323-remote-address=111.111.111.11
PortaOne-Service-Type=Voice
aggregation_key=11111
```

- The Extraction utility checks whether any records have the same aggregation key, and if so, merges these records (5).
- When the file is processed (a delay period ends or the timer fires), the Extraction aggregates the ready-to-import CDRs and adds them to the collection (6).
- The Rating utility monitors (7) the collections' storage.
- When the collection is updated, the Rating sends CDRs to the billing engine via RADIUS (8).
- PortaBilling® authorizes the user, calculates the charges for the CDRs according to defined rates, if required, and then saves the created xDRs to the main database.
- The Rating utility receives information about the import status (success, fail) (9).
- The Rating utility updates the information in the collections' storage (10).
- An administrator can control the import process via web interface, and in case of mistakes, adjust the extraction or rating settings and re-run the import (11).

xDR mediation modes

How you receive files with CDR records depends on the peculiarities of your business and your agreements with your carrier. For example, your carrier sends you:

- one file per day, or
- one file every 30 minutes, etc.

For this reason, you can import CDRs from one or several files to PortaBilling® in different modes:

- **sequential** mode (the files are imported one by one) – For example, you upload three files to an incoming folder (e.g. A, B and C). The xDR Mediator processes the CDR records from each

file and passes the CDRs to the billing engine for rating, file by file.

Here, the Extraction utility processes the CDRs from each file and creates a separate CDR collection for each of them (e.g. collection A, collection B, etc.). The Rating utility processes the CDRs from collection A and sends them to the billing server in chronological order. The Rating utility then processes the CDRs from collection B, etc.

- **consolidated** mode (the data from several files is collected and then imported) – For example, the same files, A, B, and C are in the incoming folder.

In this case, the Extraction utility creates a single collection for the CDR records from all the files. It processes the CDRs from each file and waits for 60 minutes by default. At the end of this time, the Extraction utility updates the collection and places ready-to-import CDRs to be rated. The Rating utility processes the CDRs from the collection and sends them to the billing engine in chronological order.

NOTE: Both sequential and consolidated modes enable you to import CDRs in chronological order. These modes require the use of **one** Extraction utility and **one** Rating utility. The number of simultaneous flows depends on the productivity of your system.

- **random** mode (several files are simultaneously processed) – For example, some files (e.g. A, B, C and D) are in the folder. The xDR Mediator processes the A, B, and C files at the same time and passes the CDRs to the billing engine for rating.

In this case, several Extraction and Rating utilities operate with source files. It is difficult to predict the order of data import since the size of the files and their process time may differ.

Auto-download of xDR files from external systems

The xDR mediator has built-in functionality for communicating with external systems (IMS, third-party gateways, etc.) to receive CDR source files via the FTP / SFTP protocols.

The xDR mediator can operate as either an FTP client or an FTP server. When it operates as an **FTP client**, it connects to the external system and downloads CDR source files via FTP / SFTP. To make this happen, the

administrator defines the credentials necessary for accessing the external system as well as the path to the folder from which to download the files.

When operating as an **FTP server**, the procedure is reversed. The administrator configures the entry point to the xDR mediator for the external system. The administrator creates a special user for the file transfer. The external system uses this user's credentials to access the xDR mediator and upload the CDR source files to the specified folder.

User details

General

Login:

Email:

Full Name

Password:

Confirm password:

Time Limitations

Activation Date:

Expiration Date:

Other

Role:

Time zone:

SSH Account: allow access

Save Cancel

FTP file transfer (inbound or outbound) is the recommended configuration compared to other ones (e.g. via SCP or ssh access) since it increases system security and simplifies xDR mediator support through system updates.

Retroactive billing during xDR import

To calculate charges for a billing session, PortaBilling® must consider the delay between the session's effective time – the time when the session starts / is connected, and the charging time – the time actually charged. For real-time charging this delay is very short (only a couple of seconds). When it comes to xDR import, the delay can be considerable – several hours or even days (e.g. roaming xDRs can arrive several days after actual

service usage). In the latter case, the billing periods that these charges and volume discount consumption refer to can already be closed.

Therefore, to ensure correct charging for a closed billing period, PortaBilling® retroactively bills during an xDR import and uses the tariff applicable at the time the session started to produce the xDRs. The xDRs are then added to the open billing period.

Volume discount usage is withdrawn from the closed billing period's counters.

Consider the following example:

Let's say John Doe has a "100 free roaming minutes" monthly quota with one rollover. Throughout November he makes calls for 90 minutes in roaming and on December 1 his billing period closes. Thus, 10 minutes remain from the November quota.

On December 2nd the roaming vendor sends you CDRs with John's service usage for November 30, which includes 5 minutes of calls. Upon xDR import, these charges are applied to John's November billing period and his quota counters for November are adjusted by 5 minutes. When John makes calls on December 3, his December quota now includes 105 minutes (5 minutes rolled over from November).

NOTE: If the quota has been used in the current billing period prior to xDR import for the previous one, the rollover from the previous billing period is not adjusted. Thus, if John makes a call on December 1, his rollover amount is 10 minutes.

13. Billing / invoicing

Multi-currency support

PortaBilling® allows simultaneous use of multiple currencies, including XAF. Certain groups of customers can each be managed in their desired currency. A reseller can apply his customer's currency (or multiple currencies) even when different from the one he uses to conduct a business with the operator (or upper-level reseller).

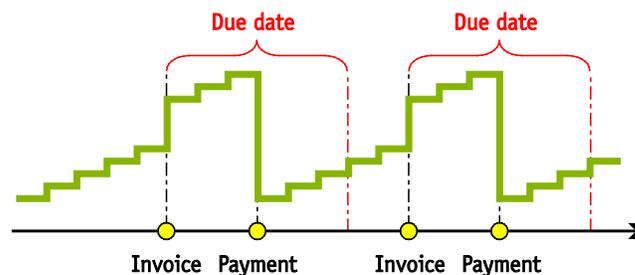
Exchange rates between currencies are either entered manually by an administrator or fetched automatically from xe.com / yahoo.com.

Obsolete currencies

In the course of time, some currencies become obsolete (e.g. coupons). The obsolete currency cannot be added to the system or selected as a base currency. But if it was added and used in the system before becoming obsolete the system will continue working with this currency as before. Note that when updating the list of exchange rates for existing currencies, if there are any exchange rates for obsolete currencies they will not be updated.

Invoicing concept

At the end of the billing period, PortaBilling® can produce an invoice for your customers. An invoice reflects all the transactions (calls, payments, refunds, subscription charges, and so on) which occurred during this period, and serves as the primary record of services provided to a specific customer, as well as his current status.

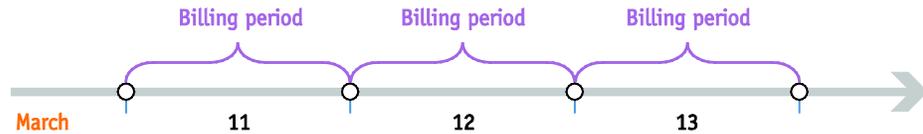


Billing periods

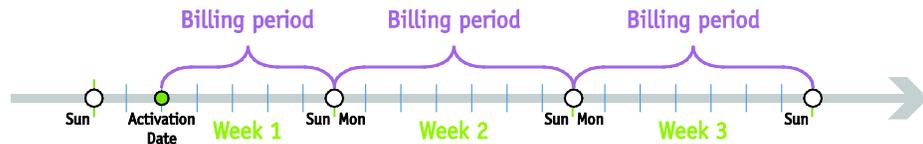
A billing period defines how often invoices / statements will be generated and when exactly they are to be produced. PortaBilling® supports the following billing periods:

- **Daily** – Covers a 24-hour period. If a customer's account is activated at 12:00 on March 11th, his first invoice covers the period

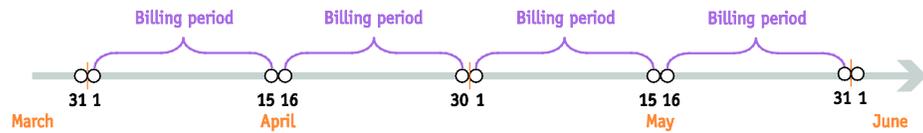
until March 12th; the second period covers the period from the beginning of March 12th until March 13th, etc.



- **Weekly** – Covers a 7-day period (Monday through Sunday). For example, if a customer’s account is created on Wednesday, his first statement covers the period from Wednesday to Sunday; the second statement covers the period from Monday to Sunday, etc.



- **Semimonthly** – Covers a period from the 1st to the 15th or from the 16th to the last day of the month.



- **Monthly** – Covers the period from the 1st of the month to the last day of that month.



- **Monthly (anniversary)** – Covers the period from the Nth day of the month to the day before the Nth day of the following month. N is the day of the month when the customer was created; therefore, if a customer was created on March 19th, their invoices will always cover the period from the 19th of the current month to the 18th of the following month.



To avoid complications for customers who were created on the 29th, 30th or 31st day of the month, their first billing period will cover the time until the 28th day of the following month, and

thereafter will always cover the period from the 28th until the 28th. For instance, if a customer was created on March 30th, his first invoice will cover the period from that day until April 28th, while his next invoice will cover the period from April 28th until May 28th.

- **30 days** – Every billing period is exactly 30 days, so if a customer was created on March 20th, his first invoice will cover the period from March 20th to April 18th, the second invoice will cover the period from April 19th to May 18th, and so on. This invoicing method allows you to make subscription fees more straightforward compared to regular monthly billing, where the same monthly fee is applied to longer (e.g. March) as well as shorter periods (e.g. February or April).



Every billing period is adjusted to the respective customer's (or vendor's) billing time zone. So, in the case of a customer with the Los Angeles time zone and a weekly billing cycle, the billing period will start at midnight on Monday Los Angeles time, while for a customer with a weekly billing cycle and the Singapore time zone it will start at midnight on Monday Singapore time. Thus while both invoices will cover 7 days (168 hours), they will actually refer to different intervals of time, with a 15-hour difference.

Closing a billing period

Or in other words: when will invoices actually appear in the system?

By default, invoices / statements are generated on the day after the last day of the billing period, e.g. invoices for customers with a weekly billing cycle are generated on Monday, while invoices for customers with a monthly billing cycle are generated on the first day of the month. A transaction is included in a certain billing period if it **was initiated** during that period; this is because calls are considered as taking place within a certain period according to their *bill time* (which is equal to their *connect time*). For example, if a call starts at 23:55 (11:55 p.m.) on March 31st and finishes at 00:43 (12:43 a.m.) on April 1st, then this call belongs to the March billing period.

For this reason, a billing period cannot be closed the next day at midnight sharp; there might be calls in progress which started just a few minutes ago, and which should still be included on the current invoice.

PortaBilling® waits a sufficient amount of time before closing a billing

period, to ensure that all calls have been completed. By default, this interval is six hours, but it can be changed via the Configuration server. Also note that since statements are generated in the billing time zone of the customer / vendor, the billing period is closed in that time zone. Thus if your system time zone is Singapore, you cannot expect to see invoices for your US customers on Monday morning, since it is still Sunday evening then in the US. Finally, in order to provide optimal system response time for your online users, PortaBilling® only performs resource-intensive calculations (such as creating statistics / invoices) during the period specified in your configuration as “off-peak”. So, in the example above, if your “statistic calculation” hours are defined as 2:00 a.m. to 7:00 a.m., you will receive invoices for US customers on Tuesday morning.

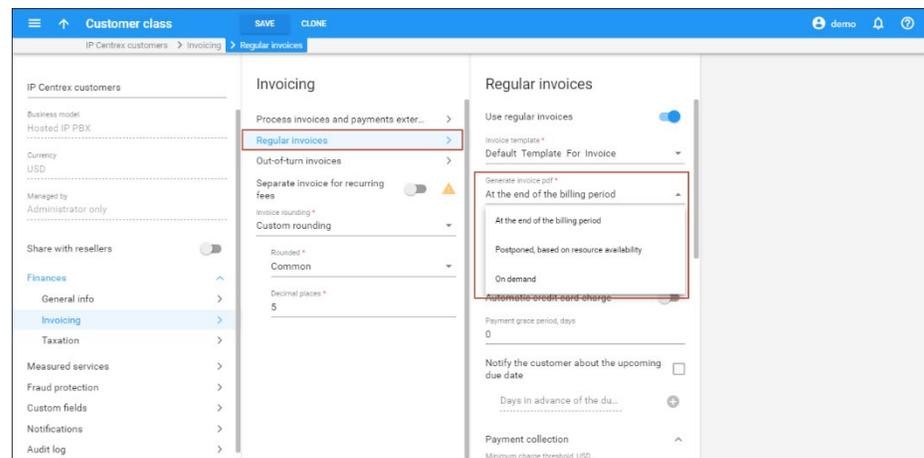
Invoice templates

You can design multiple invoice templates, so that each template has its own layout, language, logos/pictures, and the like. Every customer will be assigned a specific invoice template, according to which PortaBilling® will create a .pdf file that can be emailed to the customer and/or downloaded from the PortaBilling® web interface (by both the administrator and the customer himself). It is also possible to disable the “Generate invoice” option for a customer. In this case, PortaBilling® would only produce CDR summaries, but not generate .pdf invoice files.

Please consult the [PortaBilling Templates Guide](#) for more information about creating and managing invoice templates.

Invoice .pdf file generation

PortaBilling® offers three invoicing options configurable at the customer or customer class level:



- **.pdf file generation at the end of a billing period.** This is the default option. Once a billing period ends, PortaBilling® processes the customer's xDRs, applies charges (e.g. subscriptions, fees, etc.), creates an invoice and generates a .pdf file. XDR processing for the next customer only starts once the .pdf file for the previous customer has been generated. Therefore, it takes longer to process xDRs for all customers, though the .pdf files are quickly available.
- **Postponed .pdf file generation.** This is a useful option for service providers who automatically charge customers' credit cards. With this option enabled, PortaBilling® creates an invoice and charges a customer's credit card immediately. PortaBilling® begins to generate the .pdf files only once the calculations related to the previous billing period (e.g. xDR processing, statistics) for all customers have been completed. Postponed generation accelerates the payment procedure, evenly distributes the load on the system, allowing service providers to quickly receive revenue.
- **.pdf file generation on demand.** PortaBilling® makes all calculations for the customer, creates their invoices and saves them to a database. These invoices are then accessible via API. However, the .pdf file will not be generated unless explicitly requested. By generating .pdf file invoices on demand you can save hard disk space and decrease the workload on the system by up to 50%, depending on the data amount. Note that at any time an administrator can initiate .pdf file invoice generation on the customer's page.

PDF file generation management enables service providers to schedule invoice generation according to their business models. This results in decreasing the workload on the system, saving hard disk space and quicker statistics calculation.

Invoice parameters

Every invoice contains global invoice data, which is stored as part of the invoice record in the database:

- **Invoice number** – Unique identifier of an invoice.
- **From date** – Start date of the billing period.
- **To date** – End date of the billing period.
- **Invoice date** – Date when the invoice was generated.
- **Due date** – Date by which payment should be received.
- **Payment terms** – Description of payment terms (e.g. “due on receipt”).
- **Invoice total** – Sum of all charges in this period minus credits/refunds.

- **Invoice amount due** – Amount the customer is supposed to pay you (see below for a detailed explanation of invoice balances).
- **Invoice status** – The following statuses are possible:
 - **Open** means that the invoice has been generated, but has not yet been delivered to or viewed by the customer, so it can potentially be modified.
 - **Closed** means that the invoice has already been sent to the customer by email or downloaded by him from the web interface.
- **Invoice payment status** – The following statuses are possible :
 - **Unpaid** – The normal status after invoice generation; the customer is expected to pay the amount due (calculated as the sum of the previous invoice's amount due and this invoice's total).
 - **Overdue** – No payment was received for this invoice and the due date has already passed. If applicable for this customer class, the collection process (notifications, suspension and termination) will be launched.
 - **Paid** – Invoice paid in full.
 - **Partially paid** – A payment was received, but it does not cover the full invoice amount (*invoice total*).
 - **Do not pay** – The invoice total is zero (e.g. the customer did not use the service) or negative (e.g. the customer was given credit in an amount higher than the actual service usage) and there is no unpaid balance from the previous invoices. No payment is applicable to this invoice; it is produced only for the customer's reference.
 - **No payment required** – This invoice's amount due is lower than the collection threshold, so the customer is not required to pay it right now (but still has to pay it eventually). The collection process will not be launched if no payment is received.
 - **Previous balance remaining** – Although this invoice has a zero (or negative) invoice total, when combined with the amount due from earlier invoices there is still an outstanding balance. So the customer is still required to make a payment (to be applied to the earlier invoices), after which this invoice's status will become "Do not pay".

Invoices also contain certain data extracted from other PortaBilling® objects, which are included in the invoice's .pdf version:

- Information about the customer (invoice recipient).
- Information about the invoice issuer (your company or a reseller, if the invoice was issued on behalf of a reseller).
- Information about calls and other transactions included in the invoice.

Invoice balance

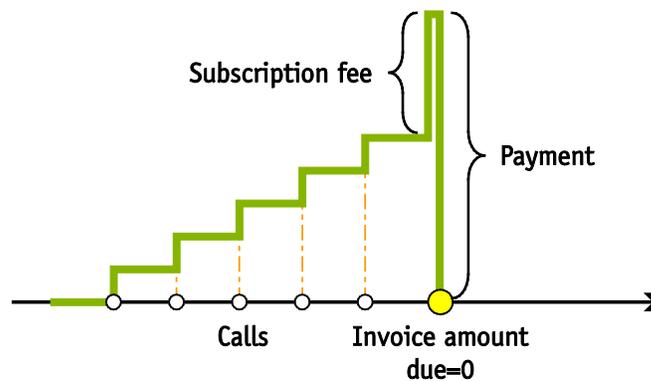
PortaBilling® provides two methods for computing the invoice balance (amount due):

- **Simple** – In this case, the invoice’s amount due is equal to the invoice total, and is calculated as the sum of all charges during the given period (no previous payments are taken into consideration). This is an optimal method for integration with an external bookkeeping system, where you keep track of your incoming payments via accounting software, and not PortaBilling®.
- **Balance-aware** – The invoice total is calculated as the sum of all charges during the given period (both call and non-call related), minus the sum of all credits/refunds. The invoice’s amount due is calculated as: $\text{previous_balance} + \text{invoice_total} - \text{payments}$. This allows you to “carry over” a balance in the case of partial payments. For example, suppose a customer’s March invoice was for \$40. If he makes a payment of \$30 on April 10th, makes calls for \$25 during April, and is also issued a \$3 refund, then his April invoice will have an invoice total of $\$25 - \$3 = \$22$, while the invoice’s amount due will be $\$40 + \$22 - \$30 = \32 .

By default, PortaBilling® uses the balance-aware invoice generation method.

Charging the invoice balance

To simplify payment collection and improve cash-flow, PortaBilling® can charge a customer’s credit card before closing an invoice. So if, as in the example above, the invoice’s amount due was initially \$32, the customer’s credit card will be charged \$32, with payment entered to the April billing period. As a result, the customer’s invoice will be created with a zero amount due.



This type of auto-payment can be activated on a per-customer basis (using the **Payment method** panel).

Payment status

Every time a payment is recorded in the system (this could be a periodic payment, an online payment by a customer, or a payment entered manually by the administrator), in addition to modifying a customer's balance it will also be applied to one of his unpaid invoices. If the amount so applied equals the invoice's amount due, the invoice becomes "paid", while if the payment is less than the amount due, the invoice becomes "partially paid". Payments are applied to an invoice cumulatively; thus if an invoice is for \$30, and the customer makes three payments of \$10, \$13 and \$17, following this last payment the invoice will be "paid". If a customer has several unpaid invoices, the payment will be applied to the oldest one.

Unallocated payments

If a payment exceeds the total amount of all unpaid invoices, the remaining sum will be assigned to the special customer property "unallocated payments", and applied to future invoices. For instance, suppose a customer receives his weekly invoice, with an amount due of \$8.99. Since he plans to leave for a three-week vacation, he sends in a payment of \$36. This entire amount is applied to the customer's balance, so that \$8.99 will cover the existing invoice, while \$27.01 will remain in "unallocated payments". When his next several invoices are created, they will show an amount due of zero and the status "paid in full".

NOTE: Unallocated payments do not represent a "cash reserve". When a payment is made, the amount is immediately applied to the customer's balance. Unallocated payments merely show that the customer "overpaid" you sometime in the past, and are used to correct the paid/unpaid status of future invoices.

NOTE: Unallocated payments are automatically applied to pay off the customer's regular and out-of-turn invoices.

You can see the customer's current "unallocated payment" status on the **Balance & Credits** tab in the customer info.

Collection policy

Customer class allows you to define a policy for automated payment collection.

Important dates

- Issue date defines the moment when an invoice was issued; all other dates are generated based upon it.
- Invoice grace period indicates the period (in days) following the invoice issue date for a payment to be received. The invoice due date is calculated as `invoice_issue_date + grace period`. For

example, if the issue date is June 1 and the grace period is 15, the invoice must be paid before June 16. The grace period can be zero, in which case the invoice is considered to be due upon receipt.

- Suspension time defines within how many days the customer will be suspended after an invoice becomes overdue. Suspension means that, while the customer's information will remain intact, he will not be able to use certain services, most notably sending or receiving calls.
- Closing time defines within how many days a customer account will be closed after an invoice becomes overdue.

Let's take a look at the following example. Customer A has a grace period of 21 days, a suspension time of 14 days, and a closing time of 21 days. A's invoice was generated on May 1st, so that the invoice due date is May 22nd. If A does not pay by that date, the invoice will become overdue; 14 days (2 weeks) after that – on June 5th – his account will be suspended; and on June 12th – 21 days after the due date, or one week after the suspension date – his account will be permanently closed.

Automated payment collection

You can automate customers' invoice payments by charging their credit cards. This saves customers' time and prevents having their service suspended since their invoices are always paid automatically and on time.

You can charge customers' credit cards either immediately upon invoice generation or on their due date.

Upon invoice generation

When a customer's billing period closes, PortaBilling® calculates their charges, creates their invoice and charges their credit card. After successful payment processing, PortaBilling® generates their .pdf invoice file with a zero amount due and the Paid in full status.

Upon invoice due date

In this case, PortaBilling® generates the invoice .pdf file and sends it to the customer. During the specified grace period, the invoice remains unpaid. On the due date, PortaBilling® charges the customer's credit card for the complete invoice amount.

A customer can pay for an invoice earlier, either partially or in full. If the payment is partial, PortaBilling® charges the customer's credit card for the remaining amount on the due date.

To have a closer look at how customers are charged on their invoice due date, consider these examples:

Regular invoices

EasyCall pays their invoices 30 days after its issue.

On February 1, EasyCall receives their January invoice with \$100 due and the due date for this invoice is March 2. EasyCall uses the service for the whole of February and receives a new invoice with \$80 due on March 1. On March 2, PortaBilling® charges their credit card for \$100 to cover the January invoice and on March 31, PortaBilling® charges them for \$80, as that is the due date for their February invoice.

Out-of-turn invoices

You can collect delayed payments for out-of-turn invoices, too. The grace period for out-of-turn invoices is 10 days. On June 10, EasyCall receives an out-of-turn invoice for their equipment rental. On June 20, PortaBilling® charges their credit card for the invoice amount.

Re-collection attempts

If an initial payment is unsuccessful for some reason, PortaBilling® can try to charge the customer's credit card again (refer to the Re-collection attempts section for details).

Let's say you have configured a re-collection attempt for 20 days after the due date. On April 1, EasyCall receives their invoice for March 1-31 with an amount due of \$100. This invoice's due date is May 1. PortaBilling® charges the customer's credit card for \$100 but the payment fails so the invoice becomes overdue. On May 1, EasyCall receives a new invoice for April 1-30 with a \$150 amount due.

Then on May 21, 20 days after the March invoice due date, PortaBilling® again charges the credit card for \$100. The payment fails again so the invoice remains overdue.

May 31 is the due date for the April invoice. The total outstanding balance is now \$250, since it includes the previous balance for the March invoice (\$100) and the April invoice's one (\$150). This time, PortaBilling® charges the credit card for \$250 and since the payment is successful, both invoices become Paid in full.

To automatically charge customers' credit cards for invoice amounts:

- Configure a payment system in PortaBilling®;
- Enable the Automatic credit card charge option for the customer class;
- Specify what the grace period is and for how many days after the invoice due date re-collection attempts will be made;

- Make sure customer's credit card details are saved in PortaBilling® and activate automatic charging for invoices.

Collection threshold

In case a customer has not actively used the service and his invoice amount is very low, (e.g. less than \$1) it does not really make sense to follow the normal collection process and request payment for that amount.

So if the amount due on a new invoice is lower than the specified threshold – no payment is yet required. If no payment is made, the balance is applied to the next invoice, etc. until the amount due on a new invoice crosses that threshold. When full payment is made for these outstanding invoices, it will be applied to all the open invoices and then they will be considered paid.

Notifications regarding payment due

Since it may happen that a customer has lost or overlooked the original invoice, PortaBilling® can send automated notifications to a customer several times before the invoice due date, stating that payment has still not been received. You can configure any set of days, e.g. “10,7,1” will send notifications 10 days, 7 days and 1 day prior to the invoice due date.

Overdue invoice notifications

When an invoice is past its due date, PortaBilling® will continue notification attempts, but in this case using a different notification text. Here again, you have complete freedom in configuring this notification policy: for instance, “0,7,14” will send an alert to the customer on the due date, and then 7 and 14 days later.

Re-collection attempts

It could be that the initial attempt to charge a customer's credit card has failed due to a temporary problem (e.g. he exceeded his credit limit, having recently made a very expensive purchase, but has now paid his credit card bill). In such a case, it would be useful to try charging his card sometime later. However, since the merchant bank will usually charge you for every failed credit card transaction, this should not be done too often.

In the customer class definition you can specify when re-collect attempts should be made, e.g. “0,3,7” means that PortaBilling® will attempt to charge the customer's credit card on the due date, 3 days after the due date, and 7 days after the due date.

Suspension warning time

As a last resort to prevent service interruption, PortaBilling® will send another alert to the customer prior to service suspension. This parameter may also be configured in the customer class definition.

Closing warning time

In the case that an invoice is not paid in full after service suspension, a reminder notification will be sent to the customer stating that his account is about to be permanently closed.

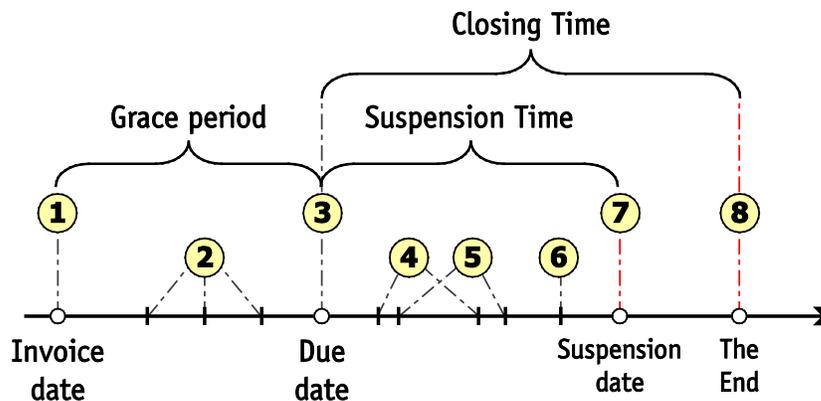
Customer account during collection

When a customer has an unpaid or overdue invoice, all call services are rendered as usual, subscription fees are charged for the current period, and new invoices are generated (for instance, if your customer has net payment terms of 60 days, he may have two other (newer) invoices generated by the time his invoice becomes overdue).

When a customer has suspended status, his call services are blocked, but some subscription fees are still applied (since his account is being provisioned, e.g. his voicemail is kept on the server).

When the customer's account is closed, no further activities, be it calls, subscription fees, self-care access, or any other, are available for his account.

Collection process



- At the end of the billing period an invoice is generated (1).
- If the invoice has a positive amount due, it is considered unpaid. When the invoice status is unpaid, a customer may be reminded of the approaching due date (2).
- On the due date, the unpaid invoice becomes **overdue** (3).

- Several collection attempts may be made (attempts to charge the customer's credit card on file for the amount due) (4).
- The customer may be reminded that the invoice is overdue (5) and that service may soon be suspended (6).
- On the suspension date, the customer's status changes to suspended (7), which automatically blocks his access to services.
- If the payment issue is not resolved, the customer's account is closed on the closing date (8).

In some cases, after the customer is suspended, and thus finally realizes that there is an unpaid invoice(!), he needs some extra time to submit payment. In this case, his suspended status may be temporary lifted by the administrator. The administrator will revert the customer's status to normal and specify the date until which suspension is postponed. If the invoice is not paid in full by that date, the customer will be automatically suspended again.

Void invoices

It sometimes happens that an error is detected after an invoice has been generated and delivered to the customer. A new invoice must be produced, but the old one must be kept for audit purposes. The void invoice operation marks the invoice as canceled (this will also be visible in the .pdf file), and then a new invoice is automatically produced.

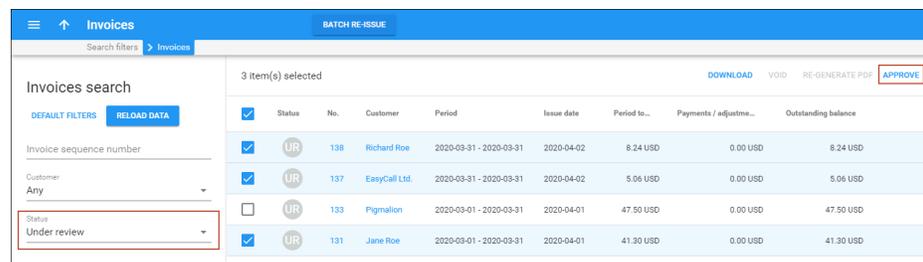
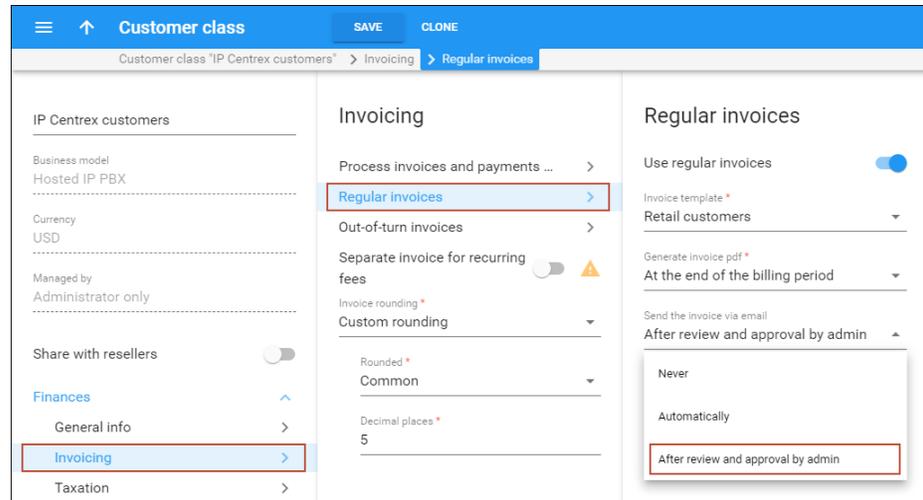
Invoice recalculation

This process voids an existing invoice and generates a new one in its place.

Invoices review

In order to provide your customers with accurate invoices, PortaBilling® allows you to review invoices before sending them to your customers. The administrator reviews the invoices and then approves them, or if necessary, performs balance adjustments (for example, to correct an invoice total) or sends them to be regenerated (if, for example, the administrator has additional reasons to update the invoice). Only then are the invoices delivered via email to your customers.

You can enable this feature for a specific customer class. At the end of the billing period, the invoices will be created but not yet sent to the customers, pending approval.



The review status of an invoice is displayed as part of the invoice record in the database. The following statuses are possible:

- **Under review** – A created invoice requires an administrator’s approval. The administrator can filter invoices by **Under review** status on the web interface.
- **Approved** – The administrator has approved the invoice and it is automatically sent to a customer via email. The invoice keeps the **Approved** status on the web interface until one of the payment statuses is applied. For instance, the **Approved** status changes to **Unpaid** or **Paid**.
- **Rejected** – If the administrator chooses to re-issue the invoice and postpones its processing to the nearest off-peak hours, the invoice status is **Rejected** until the processing is done. Once the invoice is re-issued, the status is **Under review** again, and the invoice is available for re-approval.

The invoices review functionality can be used for the following situations:

- **Billing of premium customers** – With your premium customers, you naturally want to provide them with high-quality services. Invoices are no exception. You want to be sure that all charges are included and that the invoice total is calculated accurately.
- **Billing of a new service** – An administrator creates a new service with new ratings, tariffs, products and invoice templates. If they are not sure whether all the new charges are correctly calculated

they can activate the **Send invoice after review and approval by administrator** option for the specific customer class. Only after the administrator checks and approves the invoices are they delivered to customers via email.

- **Changes in templates** – If during invoice review an administrator discovers that a template must be adjusted, they change the template. Then on the **Invoices** page the administrator selects all of the “under review” invoices that were created using this template and regenerates new ones. They are then sent for a second review.

If the “auto-charge credit card” functionality is enabled together with the “review invoices” functionality, then the credit card will be charged after the administrator’s approval of the invoice.

To prevent situations where invoice review has not been performed there is an additional field that defines how many days the invoices remain under review before they are auto-approved by PortaBilling®. Each day administrators, who belong to the corresponding mailing list, receive an email with a list of all the invoices to be reviewed. If an invoice is not reviewed by the end of the specified period, it is automatically approved and sent to the customer. An email with a list of all auto-approved invoices is sent to the administrators who belong to the corresponding mailing list.

Back-dated charges for invoices under review

When an administrator reviews invoices he can apply balance adjustments to be included on the current invoice. He specifies the amount and the date of the balance adjustment that is covered by the most recent billing period.

Consider the following example:

The administrator reviews invoice #123 for March 2015 with an invoice total of \$150, notices a human error (a \$100 subscription is charged, but it should have been free this month according to the agreement). He can click the **Refund** button and apply a \$100 credit to the customer. The invoice is then automatically scheduled for recalculation / regeneration. After that, the invoice shows a credit of \$100 and the total is changed according to this transaction. As a result, the invoice total is \$50.

Initial invoices

Sometimes you need to create a customer and immediately generate an invoice containing the activation fee, service charges for the current billing period, etc. **Initial invoice** makes it possible to apply and calculate these charges once the customer has been created. Initial invoice improves

client service by saving customers' time while providing you with a tool for generating early payments.

NOTE: A subscription can be included in the initial invoice only if its **Subscription Charges Applied** option is set to *In Advance*.

Consider the following example. The SpaceNet company signs a contract with a service provider and orders several IP phones for their employees. The service provider generates and sends an initial invoice, wherein all charges are immediately applied and calculated. SpaceNet receives this initial invoice and can now pay for all the IP phones they ordered. Thus, SpaceNet saves time by considering their future needs while the service provider benefits by immediately receiving their money.

The initial invoice tool is efficient as it optimizes the subscription process for new customers and improves the service provider's cash-flow.

Midterm invoices

It might occur that in the middle of a billing period, a customer wants to pay for services consumed and does not want to wait until the end to receive an invoice. For such customers, administrators may now generate a midterm invoice.

Let's assume that the company EasyCall, which has a monthly billing period, conducted a successful marketing campaign. Thus, EasyCall's service usage is unexpectedly high at the beginning of the month. Instead of paying a lot of money at the end of the billing period, the company wants to pay for their services consumed up to the middle of the month (May 13th), and informs the administrator about this. On May 14th, the administrator generates a midterm invoice, which covers the period from the May 1st until the May 13th. The invoice is then immediately sent to EasyCall. This way, the regular invoice that's generated at the end of the billing period will cover the rest of the period from May 14th until May 31st.

Thus, midterm invoice functionality helps administrators to manage customers who prefer receiving several smaller invoices rather than a single large one, even if the total amount is the same.

Out-of-turn invoices

Sometimes invoices for extra services such as a technician visit or an equipment purchase need to be provided at the time when such service is rendered so the customer can immediately proceed with payment. This requires an invoice that is separate from the "normal" invoice generated at

the end of the billing period. You can generate an out-of-turn invoice (to cover only a few specific items) on demand.

The screenshot shows the 'Balance adjustment' form in the PortaBilling interface. The form is divided into three steps: 1. Operation type, 2. Charge amount, and 3. Summary. In step 2, the 'Amount, USD' is set to 20.00 and the 'Service' is 'Credits / Adjustments'. The 'Manual charge for equipment' field contains the text 'Manual charge for equipment'. The 'Internal comment' field is empty. The 'Include into the out-of-turn invoice' checkbox is checked. The form has 'BACK', 'CANCEL', and 'CONTINUE' buttons at the bottom.

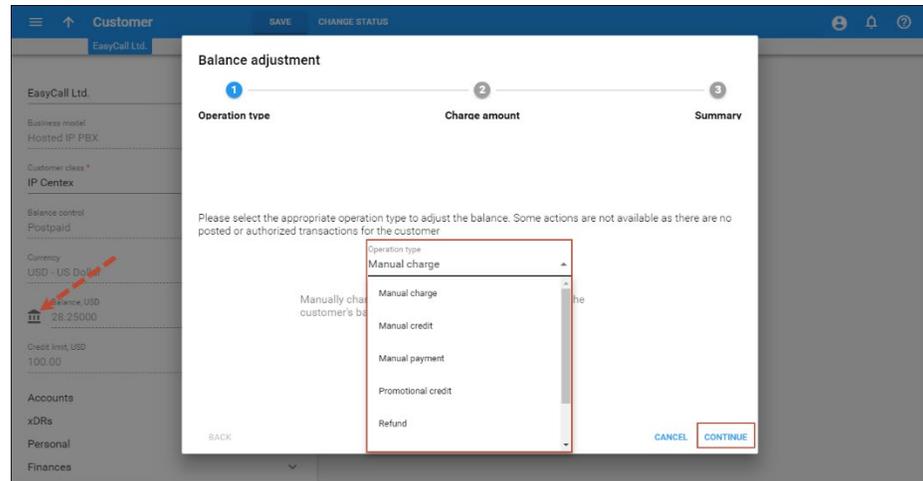
For example, let's say your customer orders 10 additional IP phones for new employees. In this case, you just issue an invoice to cover the cost of those phones and send it to the customer. Once the customer pays, you fill the order.

All you need to generate an invoice is to apply one or more “Manual Charge” transactions that contain a description(s) that will be visible for a customer and then mark it / them as **Include in Out-of-Turn Invoice**. You can then generate an out-of-turn invoice with these transactions on the **Invoice On Demand** page.

Transactions included in an out-of-turn invoice won't be reflected in the regular invoice issued at the end of the customer's billing period. Thereafter, if one or more transactions aren't included in a *generated* out-of-turn invoice during the corresponding billing period, they will be reflected in the regular invoice issued at the end of the customer's billing period.

Balance adjustment

Sometimes the administrator needs to intervene in the billing process and amend charges automatically produced by the system (xDRs for voice calls or Internet sessions, subscription charges, etc.). PortaBilling® offers a set of transactions for adjusting a customer's bill.



The currently available set of balance transaction types includes:

Transaction Type	Description	Effect On Postpaid Customer's Balance	Effect On Prepaid Customer's Available Funds
Manual Charge	Use this transaction to manually charge a customer for a specific service they used.	Increase	Decrease
Manual Credit	Use this transaction to manually give compensation related to a specific service.	Decrease	Increase
Promotional Credit	Use this transaction to give the customer credit, for example, as a sales promotion. The difference between this and "Manual Credit" is that this transaction applies to a special "Credits / Promotions" service, and not to any actual service. Basically, it provides some "virtual" funds to the customer for future use.	Decrease	Increase
Manual Payment	Use this transaction when receiving a payment (e.g. cash or check) directly from the customer.	Decrease	Increase
Refund	Use this transaction to	Increase	Decrease

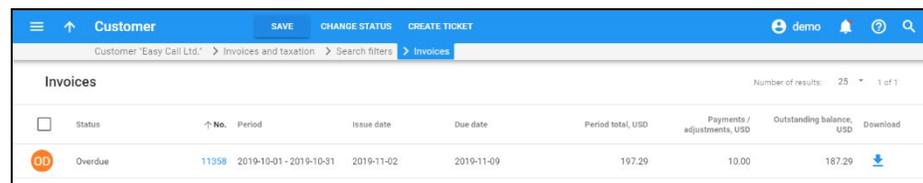
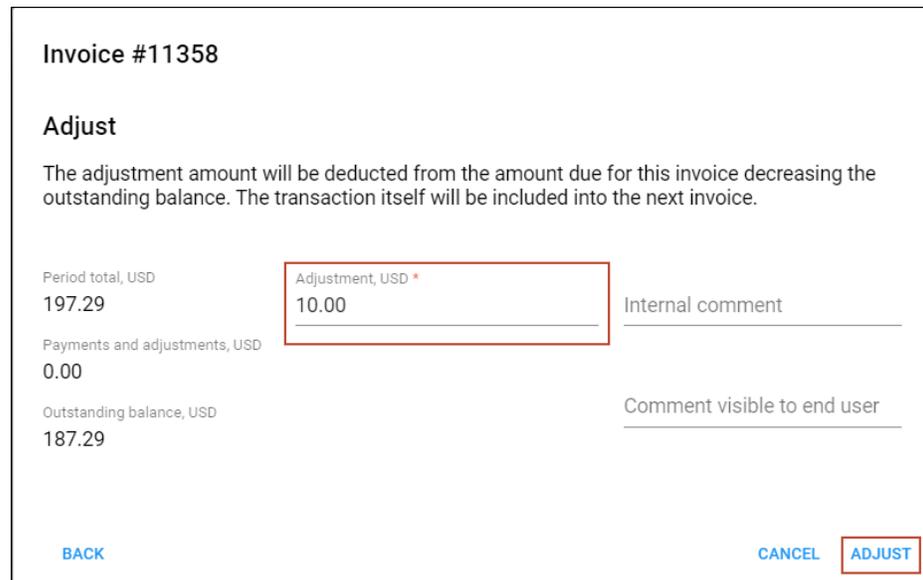
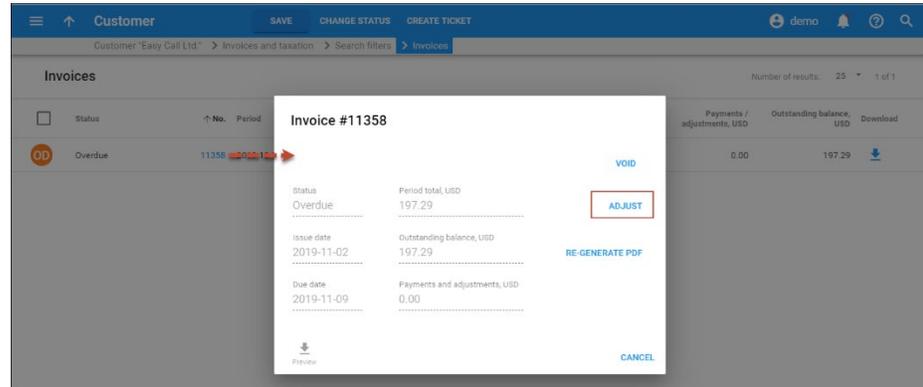
	refund an earlier payment received from the customer (e.g. a check returned by the bank).		
E-Commerce Payment	Use this transaction to charge the customer's credit card and apply the amount to the customer's balance as payment.	Decrease	Increase
E-Commerce Refund	Use this transaction to reverse a previous E-commerce payment. It withdraws funds from your company's merchant account and applies them as credit to the customer's credit card. The amount is added to the customer's balance.	Increase	Decrease
Authorization Only	Verifies that the customer's credit card is valid and reserves a given amount. Returns a transaction ID to be used in a Capture Payment transaction. Does not affect the customer's balance in PortaBilling or his credit card balance.	N/A	N/A
Capture Payment	Charges the customer's credit card and applies the amount to his balance as a payment (decreases the customer's balance). Requires a transaction ID from the Authorization Only transaction. The amount must be less than or equal to the amount of the corresponding Authorization Only transaction.	Decrease	Increase

Outstanding balance adjustment

Sometimes service providers may need to adjust an invoice amount and leave it unchanged (for example, when a customer disputes an invoice and

asks for a waiver although the invoice data has already been included in company reports).

Outstanding balance adjustment functionality does just that. Administrators can reduce the amount a customer must pay for an invoice without regenerating the invoice. To do this, the administrator just enters the waived amount for the invoice. The customer pays only the balance to cover the invoice in full.



The adjustment operation does not change any invoice data (or what's on the .pdf file). If an adjustment is made to an invoice issued several billing periods ago, all subsequent invoices are also left unchanged. The

adjustment transaction(s) is only reflected in the invoice for the current billing period.

Consider the following example. A customer's billing period is monthly. On March 1st, an invoice for February for \$1000 is generated and on April 1st, a March invoice for \$500 is generated. Later in April, the customer disputes some of the charges from his February invoice and is given a waiver for \$50. The administrator enters the waived amount and therefore the customer pays \$1450. Now his February and March invoices are paid.

These transactions will be reflected in his April invoice:

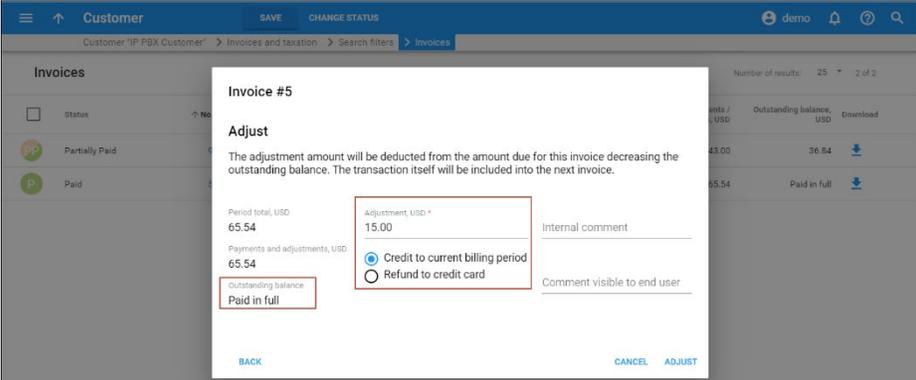
- Previous Balance – \$1500.
- Invoice Adjustment – \$50.
- Payment – \$1450.

With outstanding balance adjustment functionality, ITSPs easily deal with invoices requiring offsets. Additionally, the workload for the accounting department is reduced since activities related to invoice regeneration (i.e. resending invoices, etc.) do not have to be performed.

Adjustment for paid invoices

Administrators can also adjust invoices with a Paid status that have already been paid by customers (e.g. automatically by credit card). In this case the adjustment transaction is a refund which is applied to the customer's balance.

When adjusting a paid invoice, an administrator can select how to apply a refund to the customer: either to the customer's balance as an unallocated fund or to their credit card, if that has been configured for the customer.



The screenshot shows a web application interface for managing invoices. A modal window titled "Invoice #5" is open, displaying an "Adjust" form. The form includes the following fields and options:

- Period total, USD:** 65.54
- Payments and adjustments, USD:** 65.54
- Outstanding balance:** Paid in full
- Adjustment, USD *:** 15.00
- Internal comment:** (empty text field)
- Comment visible to end user:** (empty text field)
- Options:**
 - Credit to current billing period
 - Refund to credit card

Buttons at the bottom of the modal are "BACK", "CANCEL", and "ADJUST".

A refund that is issued for a customer's current billing period appears as an unallocated payment and is included in their forthcoming invoice.

If a refund to a credit card is unsuccessful for some reason, the refund is added to the customer's unallocated payments and is used for paying their following invoice.

This capability to adjust paid invoices improves the claims settlement process, thus increasing the customer's overall experience.

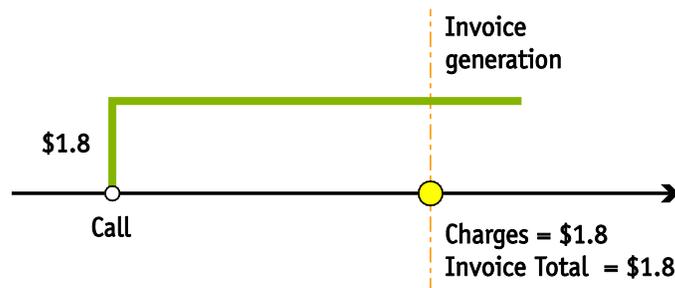
Processing taxes

There are two methods for calculating taxes: inclusive and exclusive. Inclusive means that the rate in the price list is defined with all the applicable taxes included; so then a single xDR is produced, and the total amount in the xDR includes both charges for services and the taxes applied to them. In the exclusive method, the rate is defined without any taxes, and the total amount in the xDR includes only charges for services. Taxes are calculated later on, and added as separate xDRs.

With PortaBilling® you can use both methods simultaneously to produce the same final result. In the examples given below, you will notice that, although different taxation methods are used, the information on the invoice looks the same. Deciding which method to use for a specific group of customers depends on various aspects of the two methods, which are explained in detail below.

Inclusive rate (single xDR containing charges and tax)

When you enter rates into PortaBilling®, you can define them in such a way that they incorporate the necessary charges and applicable taxes. For example, if your price is \$0.10 per minute and there is 20% VAT, the rate will be entered as \$0.12 ($\$0.10 + 20\%$), as illustrated below. When the customer is charged, the total amount in the xDR will then include the appropriate amount of taxes. In our example below, the customer makes a 15-minute call, which is charged at \$1.80 – this amount includes both service charges and taxes.



When the invoice is created, the tax information must be properly presented to the customer. Since the total amount of the invoice and the

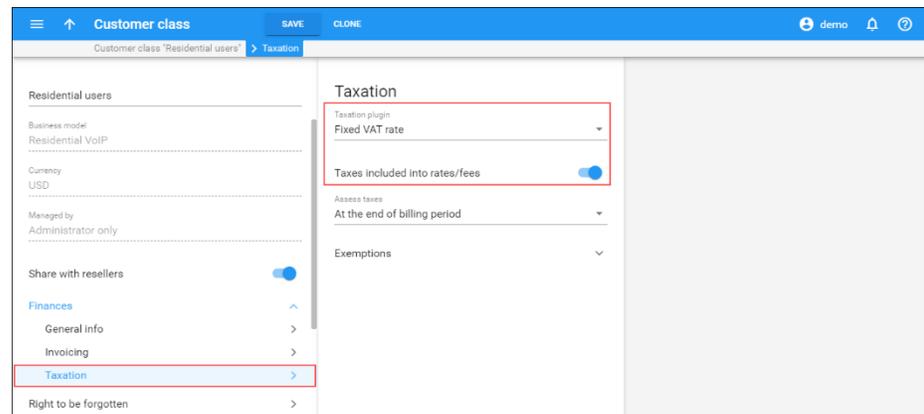
tax rate are known, the actual amount of tax and the pre-tax amount can be “back-calculated”. In our example, we assume that there was only one charge during the billing period, so that the invoice total is also \$1.8; but it works exactly the same way if there are multiple transactions in the billing period, too. By applying the tax calculation formulas, PortaBilling® determines that, since the invoice total is \$1.80 and the tax rate is 20%, the pre-tax amount was \$1.50 and the amount of tax \$0.30.



This method is ideal for prepaid services. Since every xDR produced contains the tax amount, charges and taxes are debited from the customer’s balance immediately after the service is rendered. This is also convenient for European countries, where customers are used to seeing all prices as “final.”

The inclusive rate functionality is available with these taxation plugins:

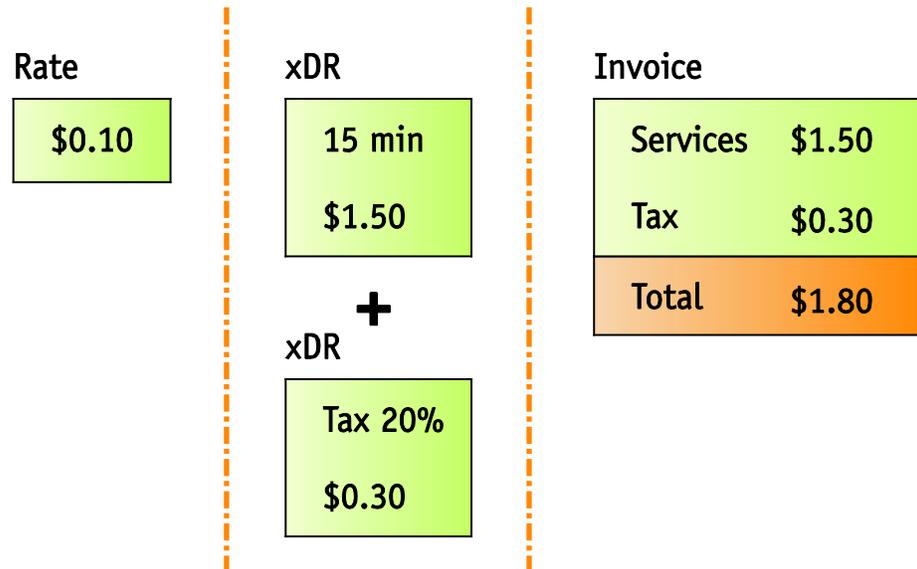
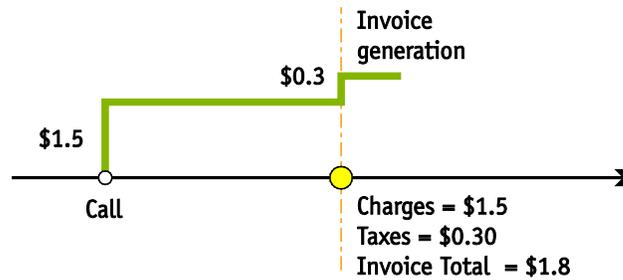
- Fixed VAT rate;
- Custom taxation.



Exclusive rate (separate xDRs for charges and taxes)

In this case, the prices you enter in the system exclude any tax information, and so when PortaBilling® processes a billing event, the charged amount recorded in the xDR is likewise exclusive of tax. Instead,

the taxation process is launched later, when closing the billing period for a given customer (this is actually the final step in closing the billing period process, after all other operations, such as charging recurring subscription fees, have been completed). All the customer's transactions during the billing period are retrieved and analyzed, and the appropriate taxes are assessed. These tax charges are added to the customer's account as separate records (xDRs), and then appear on the invoice along with the other charges.



This method allows greater flexibility when dealing with complex taxation (for instance, when the tax is not simply a percentage of the price). It is more suitable for postpaid customers, since tax amounts are applied to the balance only at the end of the billing period. It is also the preferred method for customers in the US and Canada, who are used to seeing prices without tax and having tax amounts added to a bill separately.

Since taxation rules differ from country to country, PortaBilling® supports a system of plug-in tax modules. When you configure a taxation plug-in for a customer class, you instruct PortaBilling® how to calculate and apply taxes to your customers: as either included or excluded in the rates.

Currently, PortaBilling® supports plug-ins for VAT (fixed percentage) and custom taxation, as well as the EZtax® suite, Avalara and the SureTax® taxation plug-ins. The latter three enable correct calculation of US state, country and city taxes, as well as special items such as federal USF, etc.

Unified invoice templates

Since taxes are configured by customer classes, the invoice templates in PortaBilling® are unified and apply to all taxation plug-ins. This makes invoicing configurations transparent and simplifies customer management.

Previously configured invoice templates are preserved for backward compatibility for information purposes and are available as read-only.

VAT

Fixed VAT rate is a taxation system that is used worldwide. The **Fixed VAT Rate** allows you to add a certain percentage of value-added tax. For inclusive rates, the defined percentage of tax is back-calculated from the charged amount.

Custom taxation

This taxation plug-in is similar to VAT and is specifically designed for those who need to apply different taxes depending on the type of charges (e.g. tax differently for recurring and usage charges). You can specify your own name and percentage rate value for this tax. You can apply it to:

- all services;
- all services except for subscriptions and credits / adjustments;
- only to subscriptions.

Using the EZtax suite with PortaBilling®

To use EZtax suite, you must first sign a contract with EZtax company (their tax solution is a paid subscription service; please mention that you are a PortaOne customer to receive special pricing and other promotions) and obtain the EZtax libraries for installation on your server. After EZtax suite has been enabled in the PortaBilling® configuration, the following will be done for all customers with **EZtax** assigned as the taxation method:

1. All accumulated transactions (xDRs, refunds, etc.) will be sent to EZtax suite along with the customer's information (which is used to determine his tax jurisdiction).
2. EZtax suite will calculate all applicable taxes and send them back to PortaBilling®, so that they can be inserted as extra xDRs for the given customer (each type of tax will produce a separate

record; thus if both state and city taxes are applicable, there will be two separate transactions).

3. PortaBilling® will then proceed to generate the invoice as usual.

In order to tax services properly, xDRs are mapped to EZtax transaction /service codes, which define what type of taxes apply to a given transaction. The current logic of mapping XDRs is as follows:

1. If neither the CLI (origination number or ANI) nor the CLD (destination number or DNIS) record of the call event matches the North America Numbering Plan (NANP) format (i.e. the phone number should start with 1 followed by a 3-digit area code (NPA), then a 3-digit exchange prefix (NXX), and finally a 4-digit number), then this is considered to be an international call, and the **VoIP International Usage** (19, 51) transaction/service code is used.
2. If either the CLI or CLD matches the 1-NPA-NXX-NNNN pattern, but the given NPA-NXX combination is not found in the EZtax address database (this usually means that, even though the phone number starts with 1, it belongs to some other country - for instance, the Dominican Republic), the **VoIP International Usage** (19, 51) transaction/service code is likewise used.
3. Then the location identifier for CLI and CLD are compared further:
 - If the CLI and CLD are identified as belonging to different countries, the **VoIP International Usage** (19, 51) transaction/service code is applied (calls from/to Canada fall under this category).

NOTE: Calls between the USA and Puerto Rico are defined as calls of the interstate type, and the **VoIP Interstate Usage** (19, 49) transaction/service code applies.

- Then the **state** attributes in the location information are compared for CLI and CLD. If both CLI and CLD are reported as belonging to the same state, the **VoIP Intrastate Usage** (19, 50) transaction/service code is used; otherwise the **VoIP Interstate Usage** (19, 49) code is applied.

EZtax suite updates

EZtax company provides their customers with regular updates and PortaBilling® facilitates automatic updating of EZtax suite. This provides a safe and reliable way of updating EZtax suite with no unanticipated issues emerging during tax calculation.

There are two options for the EZtax suite installation directory:

- Installation to the system's default folder.

- Installation to a separate folder for each of the system's environments.

Thus if more than one company is hosted on one common system, each company may be associated with an individual EZtax contract. This provides the ability for installing particular versions of EZtax software with corresponding updates for each separate environment. Therefore, during tax calculation the folder for the environment will be checked first, and if the folder does not exist, all of the operations will be performed within the default folder.

Tax calculation via Avalara

Avalara is one of the largest tax calculation companies in the USA. Tax calculation via Avalara is done using the XML API. Your taxation configuration is stored within Avalara and all taxes are calculated by the Avalara middleware. This significantly simplifies taxation management for your administrators.

The tax calculation process is as follows:

- PortaBilling® sends all customer xDRs, along with customer location information, to Avalara using the API.
- Avalara calculates the taxes and then returns those calculations to PortaBilling®.
- PortaBilling® uses this tax information to generate xDRs for taxes and then produce invoices.

Avalara calculates taxes based on the customers' tax jurisdiction, which is determined using PCodes – permanent jurisdiction codes. These PCodes are obtained from the customer's location (the ZIP code) and the CLI / CLD numbers in xDRs using the code mapping table provided by Avalara. This ensures accurate tax calculation.

If CLI / CLD numbers are not present in the xDRs (e.g. an xDR for subscription charges or manual credit), the customer's location PCode is used.

As a result of tax calculation, PortaBilling® receives summarized taxes grouped by tax type (e.g. federal, state, etc.) and generates separate xDRs for each tax type.

Configuration

To enable tax calculation via Avalara, first establish the account with Avalara. Contact the Avalara [support team](#) to activate the Customer Mode service. Customer Mode is mandatory to group transactions that apply to the same customer and speed up the tax calculation process.

Second, contact the PortaOne support team to configure the Avalara module on the Configuration server and install the Avalara plug-in.

Lastly, to configure the Avalara taxation parameters on the web interface:

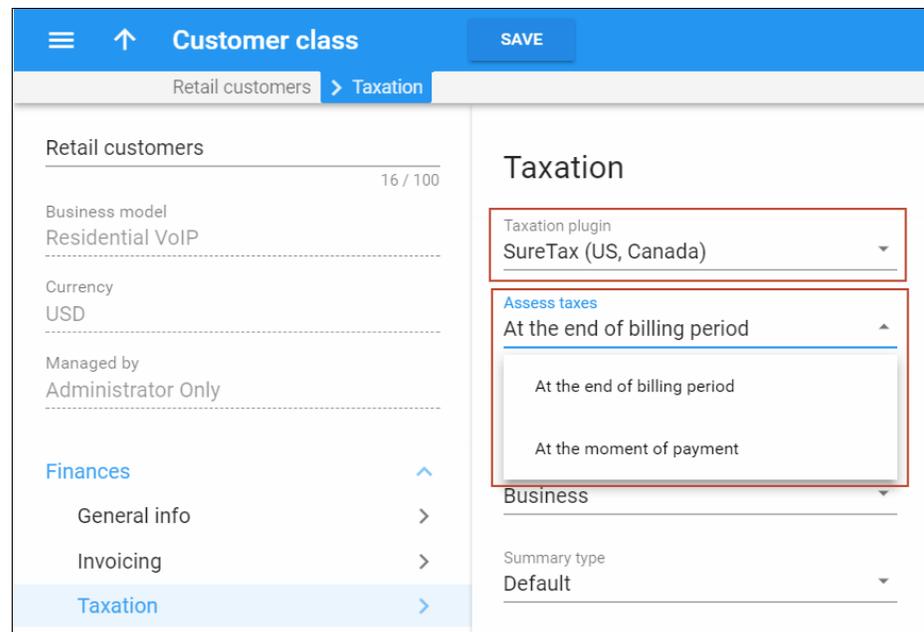
- Configure the Avalara taxation plug-in.
- Configure the taxation parameters that apply to you as the service provider on the **My Company Info** page.
- Select the Avalara taxation method and define the taxation parameters for your customers for the customer class or for an individual customer.

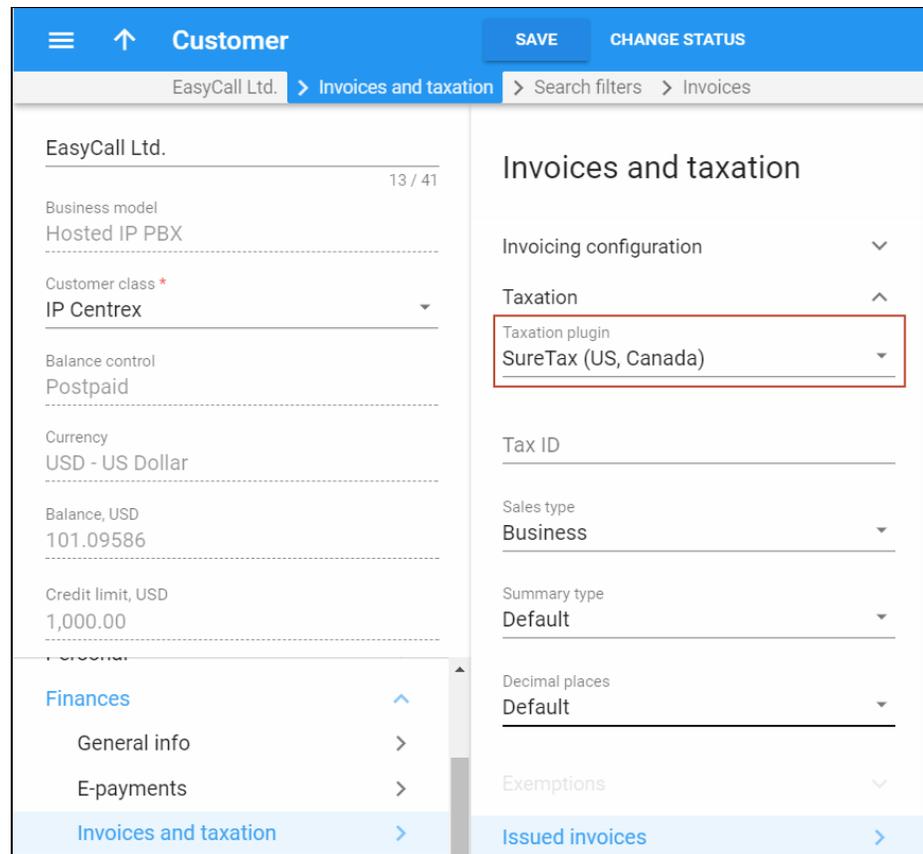
For backward compatibility, the EZTax taxation plug-in is still available. However, due to the simplified management of the Avalara updates, it is recommended that you switch to Avalara taxation.

Using the SureTax taxation plug-in with PortaBilling®

SureTax is another taxation system that allows you to accurately and easily calculate taxes for your US and Canada customers.

You can use this taxation method to tax services (e.g. voice calls, DID usage, etc.), subscriptions, bundle promotions and payments. SureTax (as a taxation method) can be defined within the customer class or for the individual customer.





The screenshot displays the 'Customer' configuration page for 'EasyCall Ltd.' in the 'Invoices and taxation' section. The left sidebar shows navigation options: 'Finances', 'General info', 'E-payments', and 'Invoices and taxation' (selected). The main content area is divided into two columns. The left column lists customer details: Business model (Hosted IP PBX), Customer class (IP Centrex), Balance control (Postpaid), Currency (USD - US Dollar), Balance, USD (101.09586), and Credit limit, USD (1,000.00). The right column is titled 'Invoices and taxation' and includes sections for 'Invoicing configuration', 'Taxation' (with 'Taxation plugin' set to 'SureTax (US, Canada)'), 'Tax ID', 'Sales type' (Business), 'Summary type' (Default), 'Decimal places' (Default), and 'Exemptions'. A 'SAVE' button and 'CHANGE STATUS' link are visible at the top right.

When you select SureTax as a taxation plug-in, you can define such parameters as:

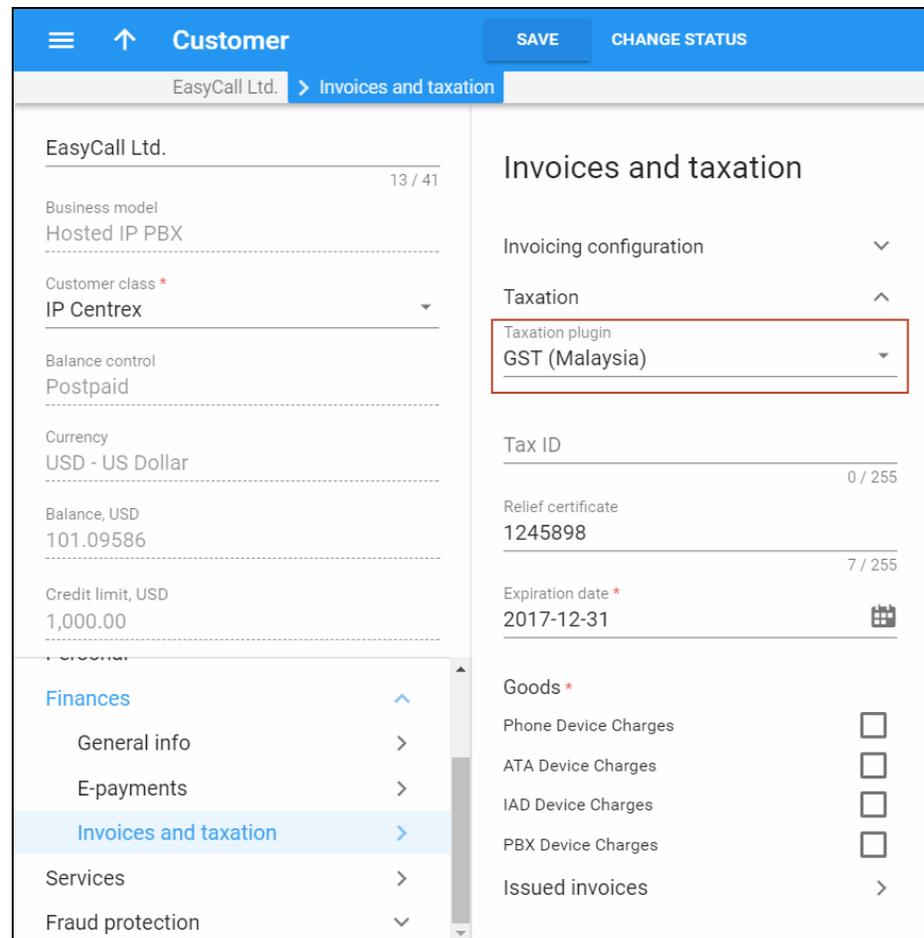
- type of customer;
- summary type;
- type of tax exemption.

Goods and services taxation plug-in for Malaysia

The **GST (Malaysia)** plug-in handles the goods and services tax (GST) that was implemented in Malaysia in April 2015.

This plug-in does the following:

- Applies 6% GST.
- Provides the option to specify tax relief certificate information, and in the case of goods that fall under the certificate, applies 0% GST.



The screenshot shows the 'Customer' configuration page for 'EasyCall Ltd.' with the 'Invoices and taxation' section selected. The left sidebar contains a navigation menu with 'Invoices and taxation' highlighted. The main content area is divided into two columns. The left column shows customer details: Business model (Hosted IP PBX), Customer class (IP Centrex), Balance control (Postpaid), Currency (USD - US Dollar), Balance, USD (101.09586), and Credit limit, USD (1,000.00). The right column shows 'Invoices and taxation' settings: Invoicing configuration, Taxation (expanded to show 'Taxation plugin' set to 'GST (Malaysia)'), Tax ID (0 / 255), Relief certificate (1245898, 7 / 255), Expiration date (2017-12-31), and Goods (Phone Device Charges, ATA Device Charges, IAD Device Charges, PBX Device Charges, Issued invoices).

- Creates individual tax xDRs for each subscription, manual charge and manual credit.
- Creates individual tax xDRs for different service subtypes, for example, voice calls to a mobile network or text messages sent to favorite numbers (service subtypes are defined based on destination ranges).
- Uses mathematical rounding on tax amounts (i.e. 1.204 rounds to 1.20, 1.205 rounds to 1.21, 1.206 rounds to 1.21, etc.).

Unified codes for tax calculation via external plug-ins

Different taxation plug-ins have individual sets of taxation codes that do not always coincide. Therefore, when a service provider switches from one taxation system to another, service management could become cumbersome and lead to misconfigurations (e.g. when an existing customer with one kind of taxation signs up for a product that uses another type of taxation).

To simplify this task, the tax processing logic has been unified as follows:

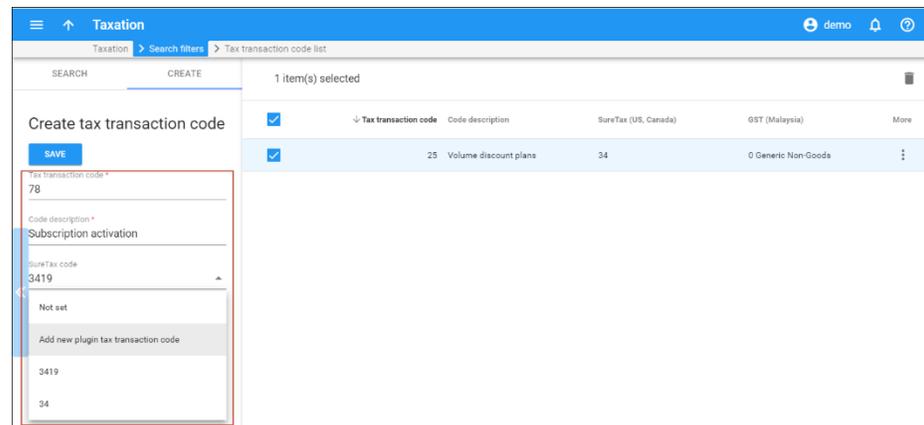
Taxation for all transactions is defined in PortaBilling® using *internal* tax transaction codes. An internal tax transaction code corresponds to the taxation code of the selected tax plug-in(s). Each set of internal tax transaction codes is unique for the administrator and their resellers.

To illustrate the tax process for customers, consider the following example:

Service provider TATA-Telecom provides voice calls, SMS and Internet access services to residential customers. They offer different promotions as part of their marketing campaign and provide users with additional products. They also have a range of IP Centrex customers and charge them for their number of active calls.

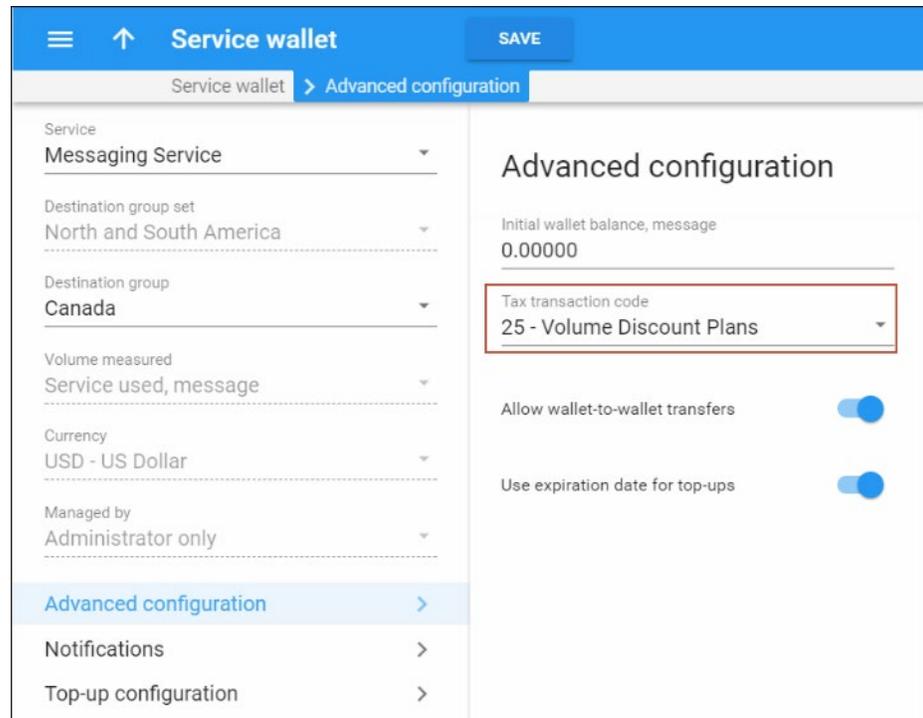
They have been calculating taxes via the taxation company A. However, they have received a better offer from taxation company B and will soon begin to use them to calculate taxes.

To simplify service management and ensure proper taxation for every customer, the administrator creates internal tax transaction codes for every taxation transaction and maps them to the taxation codes of companies A and B.



Then the administrator defines the internal tax transaction codes for each service provided:

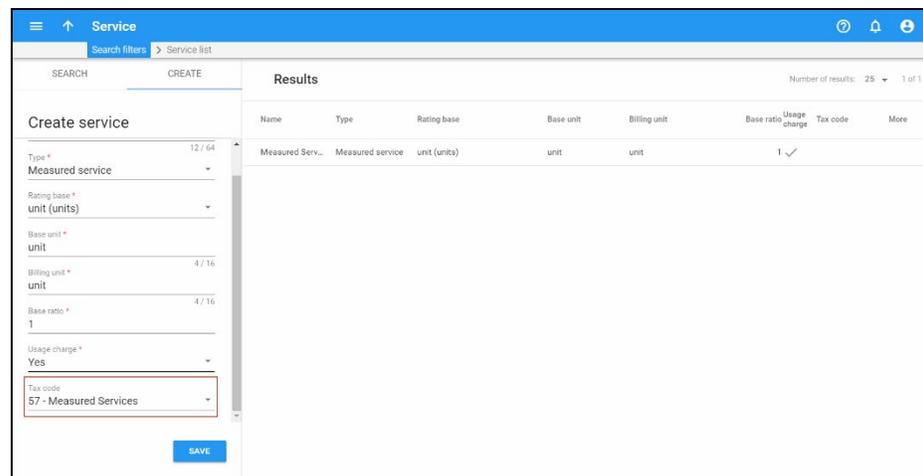
- Custom services;
- Subscriptions;
- Service wallets;
- Bundle promotions;
- DID usage.



The screenshot shows the 'Service wallet' configuration page. The left sidebar contains a list of configuration options: Service (Messaging Service), Destination group set (North and South America), Destination group (Canada), Volume measured (Service used, message), Currency (USD - US Dollar), Managed by (Administrator only), Advanced configuration (selected), Notifications, and Top-up configuration. The main content area is titled 'Advanced configuration' and includes:

- Initial wallet balance, message: 0.00000
- Tax transaction code: 25 - Volume Discount Plans (highlighted with a red box)
- Allow wallet-to-wallet transfers:
- Use expiration date for top-ups:

To apply taxes to measured services, the administrator creates a service and defines the tax transaction codes within it.

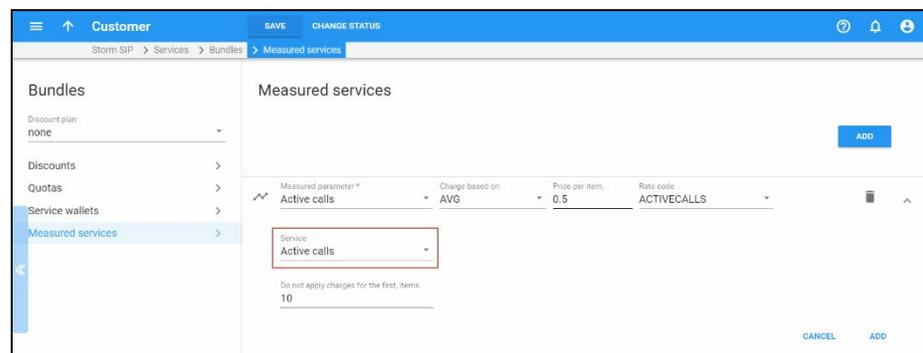


The screenshot shows the 'Service' creation page. The left sidebar has 'Create service' selected. The form fields include:

- Type: Measured service
- Rating base: unit (units)
- Base unit: unit
- Billing unit: unit
- Base ratio: 1
- Usage charge: Yes
- Tax code: 57 - Measured Services (highlighted with a red box)

 A 'Results' table on the right shows one entry:

Name	Type	Rating base	Base unit	Billing unit	Base ratio	Usage charge	Tax code	More
Measured Serv...	Measured service	unit (units)	unit	unit	1	<input checked="" type="checkbox"/>		



The screenshot shows the 'Customer' configuration page for 'Measured services'. The left sidebar has 'Measured services' selected. The main content area shows:

- Measured parameter: Active calls
- Charge based on: AVG
- Price per item: 0.5
- Rate code: ACTIVECALLS
- Service: Active calls (highlighted with a red box)
- Do not apply charges for the first, items: 10

When customers' billing periods close, PortaBilling® selects company A's taxation codes for customers that still have that configured as their taxation method and sends them to that company. Likewise, PortaBilling® selects and sends company B's taxation codes to newly added customers.

Thereby the administrator configures taxation for service wallets, subscriptions, bundle promotions and measured services in a unified manner with just a few mouse clicks.

In summary, the key assets of tax calculation using internal tax transaction codes are as follows:

1. Internal tax transaction codes must be specified for each transaction and custom service.
2. Internal tax transaction codes must be mapped to each of the taxation plug-in(s) taxation codes configured in the system.
3. Internal tax transaction codes are ignored for customers that have VAT or custom taxation.
4. Internal tax transaction codes are unique for administrators and resellers.
5. The tax transaction codes are preserved when the administrator clones volume discount plans for the same entity (e.g. when the administrator clones a discount for their own use). Consequently, once the volume discount plan owner is changed, the tax transaction code will not be cloned and must be defined later (i.e. when the administrator clones reseller A's discount and defines reseller B as its owner).

Taxation of subscriptions

A subscription can cover either a single service (e.g. an IPTV channel package) or a service bundle (e.g. 100 free minutes for domestic destinations, 200 minutes of calls to Europe and SIP phone rental for a \$20 monthly fee). In either case, subscriptions are taxed as a single unit. PortaBilling® produces a single xDR that covers the subscription's fee. PortaBilling® then sends this xDR and the tax code associated with it for taxation. Thus, if you offer several services for a subscription, the tax code must cover each of them.

For this reason, you must describe your bundle to the taxation company in detail:

- which services are included,
- the price for each service, and
- the estimated percentage of every service to be used within the bundle.

For example, you offer 100 free minutes of domestic calls for a \$30 monthly subscription. You describe your service bundle as follows:

- interstate calls: \$15 and 40% of estimated usage;
- intrastate calls: \$10 and 50% of estimated usage;
- calls to toll-free numbers: \$5 and 10% of estimated usage.

Each subscription can include up to three separate fees: an activation fee, a periodic fee and a cancellation fee. Each fee may consist of charges for one or more different services (e.g. the activation fee includes equipment costs, the periodic fee covers voice call services and the cancellation fee covers some administrative costs). To calculate taxes for such a subscription, tax codes must cover every service within a bundle and must be defined separately for each subscription fee.

The taxation company can then correctly assess each service for taxation and provide the bundle tax codes. Your administrator configures a subscription in PortaBilling® and associates the codes with it.

Depending on the taxation company, the tax codes can be of different types. SureTax defines only transaction codes while Avalara divides them into transaction and service codes.

To illustrate the taxation of subscriptions, let's say that a \$20 recurring subscription fee covers the following services:

- The SIP phone rental for \$2.
- Domestic calls for \$6.
- International calls for \$12.

The selected tax company is Avalara.

The service provider describes their service bundle to Avalara and receives the transaction / service codes. The administrator enters these codes in PortaBilling®, associates them with the subscription and assigns the subscription to the customer.

At the end of the customer's billing period, PortaBilling calculates their subscription charges and produces a single xDR record for \$20. PortaBilling® sends this xDR and the tax codes to Avalara. After the taxes are calculated and the data is forwarded back to PortaBilling®, it applies them to the customer's balance and generates the invoice.

For consistency, taxes for subscriptions and bundle promotions are calculated the same way: for a single xDR with the total amount of charges and using a tax code for the entire service bundle.

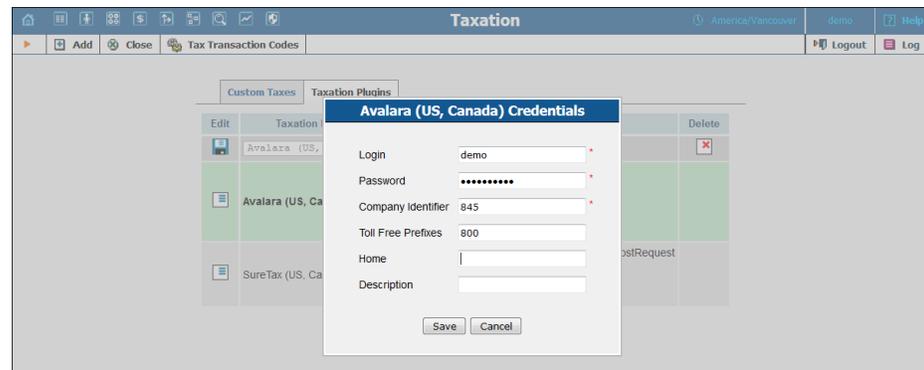
To summarize, these steps must be completed for subscription taxation:

1. Describe your service bundles to the taxation company (e.g. SureTax) to receive the bundle tax codes.
2. Configure taxation in PortaBilling®:
 - Configure the taxation system;
 - Configure the tax codes. PortaBilling® uses internal tax codes mapped to the actual tax codes. This is done to simplify taxation management for two taxation systems (e.g. when you switch from one system to another). Therefore, create internal tax codes and associate the bundle tax codes with them (Please refer to the [Unified codes for tax calculation via external plug-ins](#) section for details).
3. Configure service bundles as subscriptions and assign corresponding tax codes to them.

Taxation management by resellers

Since your resellers are separate business entities, they have different taxation parameters. To properly calculate taxes for their customers, resellers and administrators manage their taxation parameters separately.

Administrators and resellers independently configure taxation plug-ins (Avalara, SureTax or EZTax) on the **Taxation** tab and apply them to their customers.



As a result, PortaBilling® clearly distinguishes between resellers' customers and those of an administrator and ensures proper tax calculation for them.

Optionally, if you calculate taxes via SureTax, you may configure a single taxation plug-in and allow resellers to apply it to their customers. To do this, set *Yes* for the **AllowTaxPluginCredentialsInheritance** option of the Configuration server. SureTax distinguishes customer xDRs received from PortaBilling® based on the reseller ID and then creates separate tax reports for your own customers and for those of your resellers.

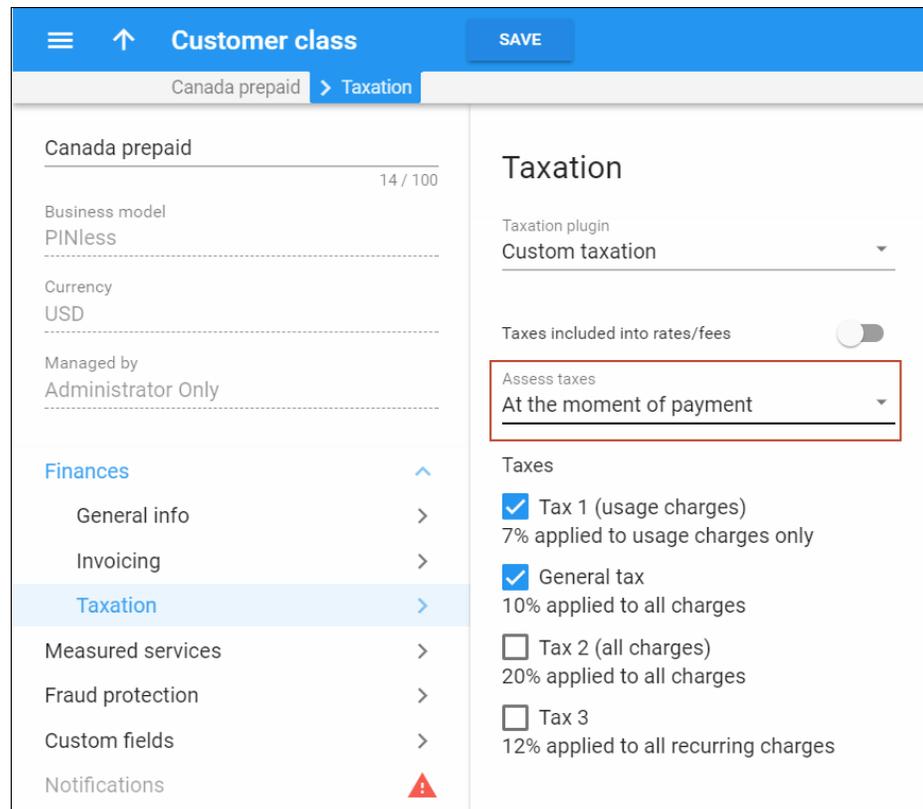
Taxation management for resellers grants additional flexibility to their service provisioning and ensures proper tax processing for their customers.

Taxes upon payment

As a general rule, taxes for consumed services are calculated at the end of each billing period and are included in the invoice that the customer receives.

However, for prepaid services such as prepaid cards, taxes are applied at the time a calling card or top-up voucher is purchased. The amount of the card excludes taxes. If later on a customer makes a payment using a credit card to top-up his / her account, the taxes must be applied at the time of purchase.

To enable taxes to be calculated and applied to customers at the time of payment, specify the tax rate and select the **Apply taxes at the moment of the payment (prepaid Customers only)** option for a particular customer class. Thus, you can differentiate among your prepaid customers and specify different tax rates for them (e.g. Canadian residential customers are charged 13% HST tax while European residential customers are charged 20% VAT tax).



The screenshot shows the 'Customer class' configuration interface for 'Canada prepaid'. The left sidebar contains a navigation menu with 'Taxation' selected. The main content area is divided into two panels: 'Canada prepaid' (left) and 'Taxation' (right). The 'Taxation' panel includes a 'Taxation plugin' dropdown set to 'Custom taxation', a 'Taxes included into rates/fees' toggle switch, and an 'Assess taxes' dropdown menu highlighted with a red box, set to 'At the moment of payment'. Below this, there is a list of taxes with checkboxes: 'Tax 1 (usage charges)' (checked, 7% applied to usage charges only), 'General tax' (checked, 10% applied to all charges), 'Tax 2 (all charges)' (unchecked, 20% applied to all charges), and 'Tax 3' (unchecked, 12% applied to all recurring charges).

This is how it works:

When a prepaid customer makes a payment via a credit card on the self-care interface the tax is calculated for them based on the amount of the top-up and is added to the payment sum. Upon successful payment processing, the customer's balance is topped up by the actual amount excluding the tax.

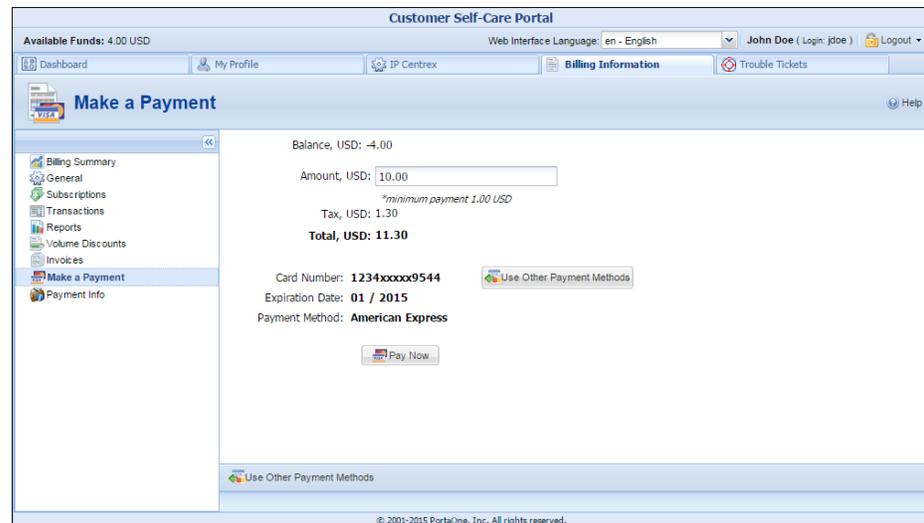
For each payment transaction, two xDRs are created and recorded in the database, one for the total payment made (the topped-up amount + tax) and another one that only contains the amount of tax.

Consider the following example:

A prepaid customer, John Doe from Vancouver, BC, buys a \$10 prepaid card and registers it for PINless dialing service.

After a while he only has \$4 in available funds so he wants to top up his account.

He visits his self-care interface and enters \$10 as the top-up amount. A \$1.30 HST tax (13%) is calculated and added to the amount entered, bringing the total sum to \$11.30.



The screenshot displays the 'Customer Self-Care Portal' interface. At the top, it shows 'Available Funds: 4.00 USD' and the user 'John Doe (Login: jdoe)'. The main section is titled 'Make a Payment'. On the left, there is a navigation menu with options like 'Billing Summary', 'General', 'Subscriptions', 'Transactions', 'Reports', 'Volume Discounts', 'Invoices', 'Make a Payment', and 'Payment Info'. The main content area shows the following details:

- Balance, USD: -4.00
- Amount, USD: 10.00 (with a note: **minimum payment 1.00 USD*)
- Tax, USD: 1.30
- Total, USD: 11.30**
- Card Number: 1234xxxxx9544
- Expiration Date: 01 / 2015
- Payment Method: American Express

Buttons for 'Use Other Payment Methods' and 'Pay Now' are visible. The footer contains the copyright notice: © 2004-2015 PortaOne, Inc. All rights reserved.

John Doe confirms the payment and his credit card is charged \$11.30. After the payment is processed, two xDRs are recorded in the database: one for the amount paid (\$11.30) and another one for the amount of tax (\$1.30). John's available funds are now \$14.

If the top-up is done using a voucher, the tax amount is not included in the voucher. However, the taxes are included in the voucher purchase (e.g. when John Doe buys a \$10 voucher, he pays \$11.30).

Thus, calculating taxes and applying them to your customers when they make payments allows you to effectively manage your tax proceeds and keep your accounting clear.

Applying taxes to auto-payments

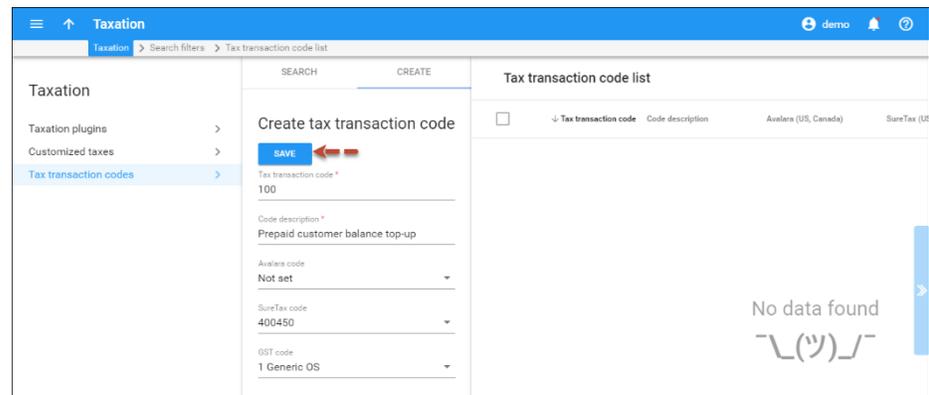
Very often, customers prefer to make regular automatic payments so that their service is uninterrupted and they do not have to bother with invoices. When customers use services on a prepaid basis, their taxes are calculated at the time of purchase, i.e. when they top up their available funds.

Typically, taxes for customers in the USA and Canada are calculated by external taxation companies such as SureTax and Avalara. When prepaid customers automatically top-up their available funds, PortaBilling® properly determines taxes and adds them to the total sum of amount charged.

Since a prepaid amount can cover multiple services and / or destinations which might be taxed differently, it is necessary for the ITSP to estimate the percentage of services their customers will use (e.g. for local, long distance and international calls and domestic SMSs) to know which taxes to apply. To do that, the ITSP contacts the taxation company (e.g. SureTax) and describes their service needs.

Based on that data, SureTax creates a taxation bundle for the ITSP and provides a bundle ID (e.g. 100). This bundle ID is then used as a tax transaction code for a main balance top-up in PortaBilling®.

To configure PortaBilling® to calculate and add tax amounts to prepaid customers' auto-payments, an administrator must map the bundle ID received from the taxation company to the internal transaction code.



Then the administrator creates a customer class and specifies when to assess taxes, defines the tax transaction code and assigns that customer class to a prepaid customer.

The screenshot shows the 'Customer class' configuration page for 'Prepaid customers' in the 'Taxation' tab. The page is divided into two main sections: 'Prepaid customers' on the left and 'Taxation' on the right. The 'Prepaid customers' section includes fields for Business model (Residential VoIP), Currency (USD), Managed by (Administrator only), and Share with resellers (toggle). The 'Taxation' section includes fields for Taxation plugin (SureTax (US, Canada)), Assess taxes (At the moment of payment), Tax transaction codes (12 - voice calls), Regulatory code (VOIP), Sales type (Residential), Summary type (No summary), and Decimal places (2). A red box highlights the 'Assess taxes' and 'Tax transaction codes' fields.

For example, let's say your taxation company is SureTax and your prepaid customer, "ABC," wants to automatically make a payment of \$90.00 each time their available funds drop to \$10.00.

When the auto-payment is initiated, PortaBilling® sends a request to SureTax with the tax transaction code and the amount to be paid (\$90). SureTax calculates the taxes and PortaBilling® receives \$6.3 in taxes from SureTax and then adds it to the payment amount. The sum of \$96.3 is then withdrawn from the customer's credit card.

When the transaction is processed, PortaBilling® creates two xDRs: one with the amount charged (\$96.3) and the other with the tax amount applied (\$6.3).

NOTE: The number of xDRs may depend on the number of taxes applied.

The administrator can also override the existing taxation parameters for an individual customer on the **Edit Customer** page.

The screenshot shows the 'Customer' configuration page for 'Richard Roe' in the 'Invoices and taxation' section. The left sidebar lists various configuration areas, with 'Invoices and taxation' highlighted. The main content area includes fields for 'Currency' (USD - US Dollar), 'Available funds, USD' (204,470.00), and 'Taxation plugin' (SureTax (US, Canada)). The 'Assess taxes' dropdown is set to 'At the moment of payment', and the 'Tax transaction code' is set to '12 - voice calls'. Other options include 'Sales type' (Residential) and 'Summary type' (Default).

NOTE: If you override the taxation parameters for a customer, you must define the main balance top-up tax transaction code.

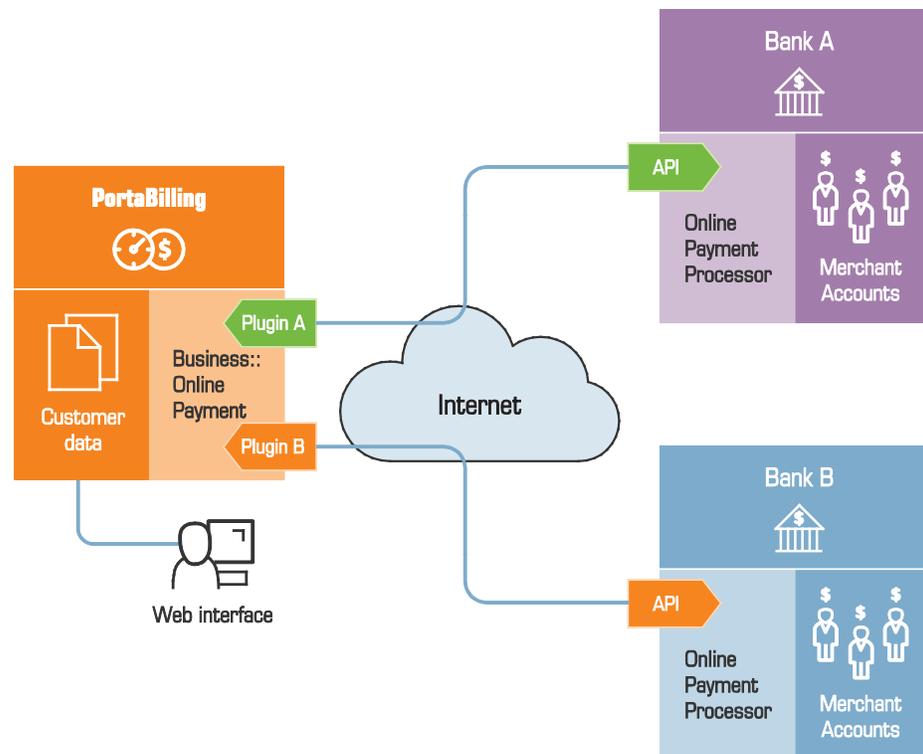
Payments

PortaBilling® allows payments to be processed online, without the intervention of an operator, by charging the customer's credit card or debiting his bank account. Payments may be initiated by:

- Your subscribers from the web interface (called “online payments” because payment is triggered by the customer, who is using the web interface at the moment the payment is made).
- The PortaBilling® system itself. These payments are called auto-payments, because they usually happen automatically.

Payment flow

The following picture illustrates the individual components of the online payment system:



Merchant account

A merchant account must be opened at a bank. Its purpose is similar to that of a normal checking account. It stores funds you have received from credit card payments.

Online payment processor

A merchant account is usually enough to start processing credit card payments via charge slips. However, if you wish to initiate transactions from your own server via the Internet, this service is provided by online payment processors such as Authorize.Net. You establish an account with an online payment processor, providing it with your merchant account information, and in return you receive credentials (username, password, etc.) for using its API. So now your application can connect to the API's server and, upon providing valid authentication information, initiate transactions.

Business::OnlinePayment framework

Unfortunately, there are many different online payment processors, most of them using their own API which is different from that of other providers. So it seems that every e-commerce application such as PortaBilling® must be programmed to support them all, which is an extremely laborious task. Fortunately, PortaBilling® provides a solution. Business::OnlinePayment is a framework which provides the necessary

middleware between e-commerce applications and online payment processors.

- The application is provided with unified API methods, e.g. `doTransaction`, so the same code in the application can be used to communicate with any online payment processor.
- A communication plug-in is created for each online payment processor. This plug-in knows such specific parameters as the server's address, the correct way of encoding transaction parameters, the proper way of parsing the server's response, and so on.

On the application side, therefore, only a few things are required: loading the correct plug-in and invoking a transaction method. After that control is handed over to the plug-in, which handles all communication with the remote server.

The whole `Business::OnlinePayment` project is open-source, and is maintained by a large team of individuals. At the moment more than 20 different plug-in modules are supplied with `PortaBilling®`, and it is easy to create new plug-ins for online payment processors. You can find the list of supported online payment processors in **APPENDIX F. Supported Online Payment processors**.

PortaBilling® payment support

`PortaBilling®` stores required information, such as customer name and address, credit card number, and so on, in the database. Multiple merchant accounts are supported (for instance, your merchant account A accepts payments in USD, while account B takes payments in euros). You may define a “minimum allowed payment” and a list of supported card types (VISA, MC, etc.) for every merchant account.

When payment is initiated (either by the customer from web self-care, or automatically) the system connects to the online payment processor and performs the transaction. The online payment processor verifies the credit card information (and performs optional steps such as address verification or CVV control). Upon receiving confirmation that the transaction was successful, `PortaBilling®` writes transaction information to the `PortaBilling®` database and modifies the balance for the account / customer.

Depending on the type of payment processor being used, both the time when `PortaBilling®` sends the request and receives transaction confirmation will differ. Thus, two separate events are performed within a single e-commerce payment transaction and are reflected in an xDR record:

- **Transaction date** – is the date upon which a customer makes a credit card payment and PortaBilling® sends a request to the payment processor.
- **Posting date** – is the date upon which PortaBilling® receives the transaction confirmation from the payment system and creates an xDR record. Thus, the posting date coincides with the date of the xDR record.

The screenshot displays the 'xDR History for Customer 'Richard Roe'' interface. It includes a summary of transactions and a detailed table of DID Usage and Payments. A red box highlights the Date/Time column in the Payments table, showing a transaction date of 2017-03-21 09:00:56 and a posting date of 2017-03-21 09:00:55.

Customer		Richard Roe	
From Date	2017-03-20 09:01:11	Total Transactions	3
To Date	2017-03-21 09:01:15	Credits	0.00000 USD
Show xDRs For	Customer and Credit Accounts	Charges	0.00000 USD
Service (s)	All	Recurring Charges	11.00000 USD
		Recurring Credits	0.00000 USD
		Usage Charges	0.00000 USD
		Usage Credits	0.00000 USD
		Payments	-53.50000 USD
		Refunds	0.00000 USD
		Taxes	3.50000 USD
		Total	-39.00000 USD

DID Usage						
DID Number	Description	From Date/Time	To Date/Time	Amount, USD	Hidden	Revert
16045566637, CA, British Columbia	DID Number Activation Fee	2017-03-21 05:02:56	2017-03-21 05:02:56	11.00000		
Subtotal				11.00000		

Payments							
User	Account	Description	Comment	Date/Time	Amount, USD	Hidden	Revert
		Payment/Refund transaction	E-Commerce Payment	2017-03-21 09:00:56	-53.50000		
Subtotal							

NOTE: In most cases, the posting date comes after the transaction date. However, the posting date comes before the transaction date in cases when postpaid customers prefer to pay their invoices automatically by credit card.

Supported payment methods

The choice of available payment methods depends on the online payment processor. In general, however, PortaBilling® supports the following payment methods:

- American Express
- E-check (direct debit from bank account)
- Direct Debit
- Discover
- Switch / Maestro
- MasterCard
- VISA.

Service providers and their resellers can also accept Bitcoin crypto currency payments for services provided. Refer to the [Bitcoin Support in PortaBilling®](#) section for more details.

Recurring payments without storing credit card info

Auto-payment helps improve the payment process for both the end user and the service provider since it is convenient and requires minimum effort on either part. When configured by either an administrator or end user, recurring payments are automatically charged to the user's credit card and applied to the user's balance.

However, auto-payment requires that the service provider store the end user's credit card information. Typically, this means that the service provider must obtain PCI DSS or similar security certification, a time-consuming process that requires periodic updating.

There is an alternative method for dealing with this situation. It takes advantage of the fact that some payment processors can store credit card information securely on their servers. They do this by generating a token after the initial transaction. The token is not the credit card information but an identifier of it, and so the token can be used for subsequent payments.

With token-based payment processing, the payment process proceeds as follows. When a customer makes a first-time online payment via the self-care interface, the credit card information entered is stored for use with future payments and forwarded to the payment processor. If the transaction is successful, the payment processor returns a token, which is stored, and when the next payment needs to be made, it is made automatically based on the stored token rather than the actual credit card data.

PortaBilling® supports token-based payment processing for the Secure Hosting and Payments, Authorize.Net, CyberSource, Payment Express and Virtual Card Services payment processors. Future releases will support additional token-based payment processors.

With token-based payment processing Auto-payment can be launched almost immediately since no PCI DSS certification is required.

Payments via external payment systems

Along with payment methods such as paying via credit card, there are online payment systems that provide users with a virtual wallet (e.g.

PayPal). PortaBilling® allows your customers to pay via such payment systems directly from their personal computers within web browser.

The process works like this: a customer logs into his customer self-care interface, chooses an appropriate payment method, clicks the **Pay** button and is redirected to a third-party website for processing. Once the user confirms his payment information he is redirected back to his account in PortaBilling® where he can view his balance update.

NOTE: External payment systems do not permit recurring (automated) payments.

PayPal

Let's take a practical example to better understand the payment flow process. Let's assume that the customer uses PayPal to pay for consumed services. The customer logs into the customer self-care portal, chooses PayPal, clicks the **Pay Now** button and is redirected to the PayPal site where he logs into his PayPal account, chooses the funding source and confirms his payment by clicking the **Pay** button. PayPal then processes the payment and sends the payment information to PortaBilling® where it is recorded and the customer's balance is updated. The customer is then redirected to the self-care portal where he can view the payment status.

Payarena

Payarena is the payment system in Nigeria using which Nigerian service providers can offer their customers extensive possibilities for payment operations plus manage their e-commerce payments in PortaBilling®.

The payment flow is similar to that of other external payment processors. When a customer initiates a payment from their self-care interface, they are redirected to the Payarena portal to provide their credit card information and confirm the payment. Upon payment processing the customer is redirected back to the self-care interface where they see their payment status and newly modified balance.

Remita

Remita is another external payment system in Nigeria that expands the list of payment processors available for managing e-commerce payments in PortaBilling®.

The payment flow is very simple. When a customer initiates a payment on their self-care portal, PortaBilling® redirects them to the Remita portal where they can fill in their credit card details and submit payment. Once payment is processed, the customer is redirected back to their self-care portal in PortaBilling® where they can check their payment status and their updated balance.

All sensitive information such as credit card details, CVV code, etc. are handled by Remita. Thus, service providers avoid system scrutiny for passing PCI compliance. This enables them to launch their e-payment processing almost immediately.

Implementation specifics:

- Remita processes payments starting from 50 NGN. Thus, an administrator must define 50 NGN as the minimum payment for the Remita payment system on the **Currencies** panel.
- To process a payment, a customer e-invoice is generated on the Remita portal. The payer's phone number and email address are obligatory for the e-invoice. These fields are prepopulated with the data pulled from the end user's account in PortaBilling®. Therefore, the phone number and email address fields on customer / account records must be filled in.

PayNearMe

PayNearMe is yet another payment system that allows your customers to pay for services in cash at any 7-Eleven or ACE Cash Express store across the United States. This option provides great value for users who may not have access to a credit card or bank account.

The payment process is quite simple. A customer views the PayNearMe payment slip on his self-care interface and prints it. The payment slip contains the bar code to be scanned at a 7-Eleven or ACE Cash Express location. Customer pays cash to the store clerk, who then records the transaction – the payment information is immediately delivered to the service provider via the PayNearMe network. PortaBilling® records the transaction and updates the customer's balance so the customer can immediately start using the service.

Receiving payments from external payment gateways

Yet another alternative to paying via credit card is to pay in cash. The payment can be made via an external payment system. The information about the transaction is then sent to PortaBilling® via the external payment gateway and PortaBilling® records the transaction and updates the customer's balance.

One example is receiving payments made via payment terminals. This payment method is one of the most popular and frequently used in the Russian Federation and other developing countries. The payment process is the following: the customer chooses a service provider on the screen of a payment terminal, enters his customer ID or account ID, inserts cash

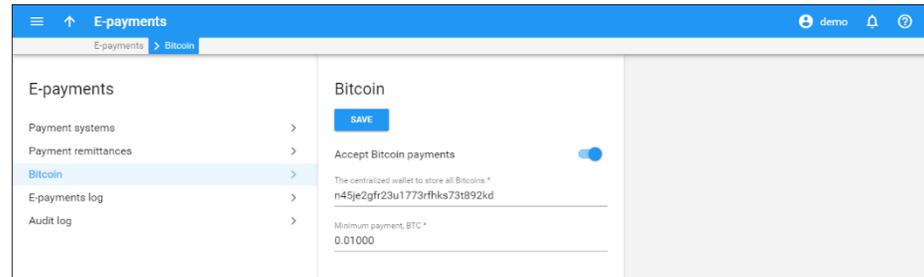
into the payment terminal and confirms the payment. The payment is processed and a request is sent to PortaBilling® along with confirmation of successful payment. The request is then processed by PortaBilling® and the customer's balance is updated.

Some minor integration with PortaBilling® is required for supporting a specific external payment gateway. This is performed with the help of a plug-in created by PortaOne.

Bitcoin support in PortaBilling®

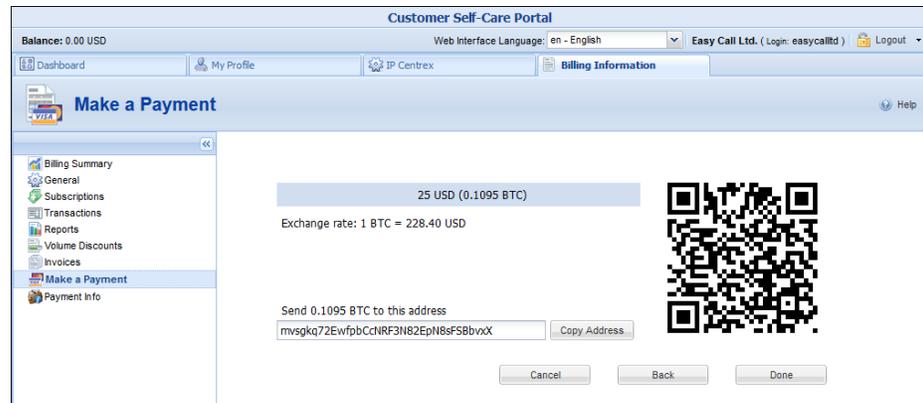
Service providers and their resellers can accept Bitcoin crypto currency payments for services provided.

To enable Bitcoin payments in PortaBilling® an administrator / reseller needs to specify the address of a centralized Bitcoin wallet which will store all received bitcoins and define a minimum payment.



For each end user, the centralized Bitcoin wallet generates a unique address where payments are sent. This address is shown on the self-care interface when the end user chooses the **Pay with Bitcoin** option. The end user then sends the top-up amount to this address from his / her individual Bitcoin wallet. Once the payment is received, it is transferred to the centralized Bitcoin wallet that is owned by the service provider or reseller and that end user's balance is topped up.

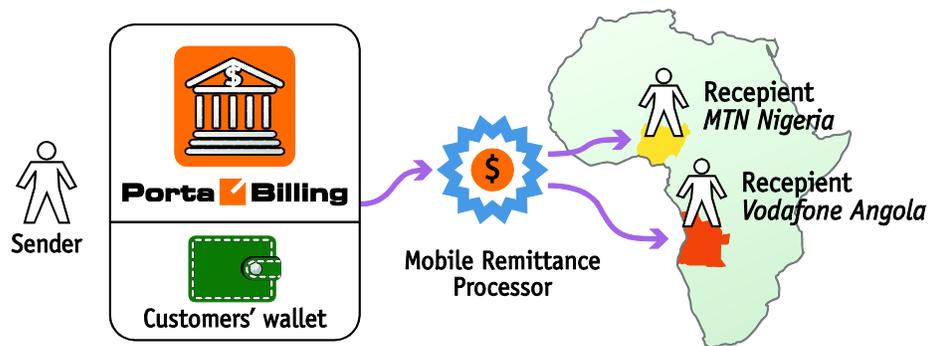
For user convenience, the Bitcoin address can be scanned with a QR code.



The principal strength of Bitcoin is that it offers anonymity to end users. When a Bitcoin transaction takes place, no personal identifying information such as name and address are required. With Bitcoin support in PortaBilling®, service providers and resellers are able to offer their services to their customers who care about privacy and want it preserved.

Mobile payment remittance

This service allows your customers to transfer funds from their accounts in PortaSwitch to a mobile phone’s balance in another country. This service is very popular in expat communities, since people who regularly use a VoIP service to call their families back home can also top up the balance on their relative’s cell phone.



What is a remittance processor?

A remittance processor (such as TransferTo) is a transaction processing system that delivers international top-up services. You create an account with an online remittance processor, provide it with your company information and establish the method of transferring funds from you to it on a periodic basis. In return, you receive credentials (username, password, etc.) for initiating transactions via the remittance API. Now your application can connect to the API server and, upon providing valid authentication information, request that funds be transferred to a customer of a mobile carrier.

The system functions in the following way:

1. Customer A subscribes to VoIP services provided by company X using the PortaSwitch platform. A adds funds to his account in PortaSwitch (cash payment, credit card, voucher, etc.).
2. A wants to send funds to the mobile phone of his relative, B.
3. A uses PortaSwitch (via web interface, IVR or SMS) to initiate a mobile remittance transaction. Funds are withdrawn from his account in PortaSwitch and the remittance processor Y is instructed to transfer the funds to B's mobile phone.
4. Y contacts the mobile carrier that owns the network that B's mobile phone is connected to, and instructs that funds be added to B's account.
5. Later (on a periodic basis), service provider X transfers funds to the remittance processor to cover the total of all such transactions in the period.

Relationship between transferred and debited amounts

By the nature of international mobile balance transfers, the recipient of the payment will receive the funds in a currency different from the one used by the “donor.” Hence the question of exchange rates arises. The service provider naturally wants to ensure that:

- The amount that is withdrawn from the donor's account is no less than the actual cost (the amount that the service provider will have to pay to the remittance processor to cover the funds actually delivered to the recipient, given the exchange rate and fees applied by the remittance processor).
- The amount that is withdrawn from the donor's account incorporates a profit margin for you as the service provider.

For example, when a transfer of 50 Australian dollars is made, if the amount payable to the remittance processor (given the AUD:USD exchange rate² and the remittance processor's transaction fee) is 52.369 US dollars, then the amount debited from the donor's account will be \$53.37, since it includes your profit of US \$1.00.

- The amount that is withdrawn from the donor's account is represented by a number that is easy for the end user to understand. Payment remittance cannot perform top-ups with an arbitrary amount (e.g. you cannot request that an amount like 0.05 AUD or 9.99 AUD be added to the remote mobile phone's balance), and so usually there is a predefined set of allowed values for each mobile network.

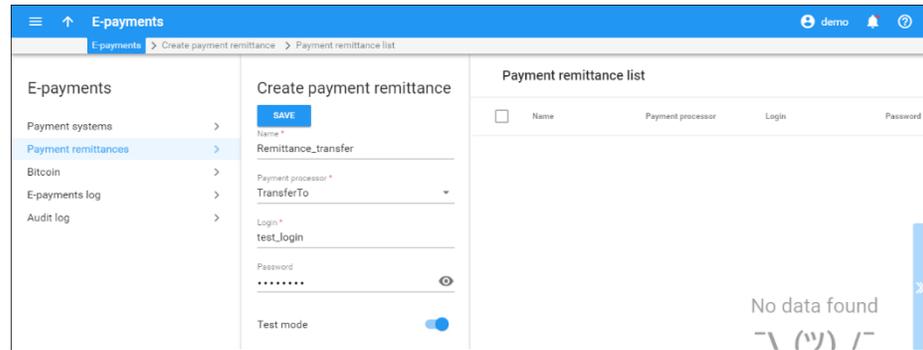
² At the time of writing

This list depends not only on the target currency, but sometimes also on the details of the arrangement with a given mobile operator, e.g. for the Aircel network in India the allowable values might be 500, 1000 and 5000 rupees, while for the Idea! mobile network (also in India) they will be 200, 400, 800 and 2000. So if a chosen “target” amount is converted into the “debit” amount using only the exchange rate, it will most likely result in an awkward number with decimal digits. Clearly, in most situations you will want to announce a nice “round” number to the user, i.e. the IVR announcement, “Your account will be debited fifty-four dollars” is better than “Your account will be debited fifty-three point thirty-seven dollars.”

In order to address all of these challenges, the PortaSwitch administrator configures a special “mapping” tariff that flexibly defines the amounts allowed to be added to a remote user’s balance and the corresponding amount to be debited from a PortaSwitch account. For instance, in order to allow customers with balances in USD to transfer 2,000 Indian rupees to a remote user on a mobile network, the following will be done:

1. The administrator creates destination INR2000.
2. Then he creates a special tariff (in USD currency and using the Messaging service) that will be used to authorize mobile remittance operations for customers with balances in USD.
3. Finally, this tariff’s administrator will create a rate for destination INR2000 and specify the Unit Price as the amount that should be debited from the donor’s account – in this case, 2,000 Indian rupees is about US \$39, so taking into consideration exchange rate fluctuations, remittance processor markups and his own profit, the administrator will set the value at \$45.
4. By setting only certain “allowable” transfer values, the administrator can ensure, for instance, that transactions are not too small (where overhead is too high) or too large (where risk of fraud is too high), even if such too small/large top-up amounts are offered by the remittance processor.
5. When a customer attempts to initiate a transfer to a phone number (9111111111), PortaSwitch queries the payment remittance server as to whether a transfer to this number is possible (it could be that the number is in a country where the remittance service is not yet available, or it is not a mobile phone number, but a landline).
6. Then a list of allowable top-up amounts is retrieved and for each of them the authorization tariff is checked; the customer will only be offered the amounts that are listed.
7. The customer selects the desired top-up amount, and after he approves the announced amount to be debited PortaSwitch sends a request to the remittance processor to initiate the top-up and debits the amount from the customer’s account.

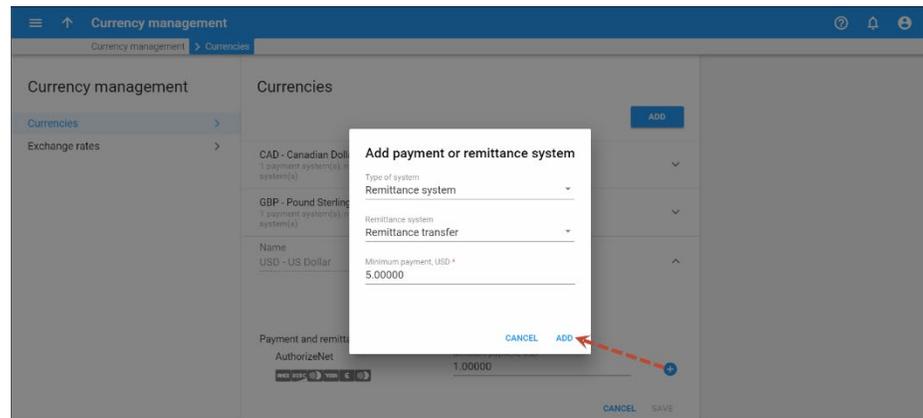
Payment remittance system configuration



1. Create a payment remittance system with the TransferTo processor on the **Payments** page.

To add a new payment remittance system, click **+ Add** and enter the account name, login, password and processor in the edit row at the top of the listing.

Test Mode is a special mode of interaction with the system. It is useful during the initial setup phase, when you may want to test the setup without actually transferring funds.



2. Define the appropriate currency to the payment remittance system on the **Currencies** page.

Click **+ Add** and fill out the relevant fields. Select the corresponding payment remittance system.

NOTE: Only those customers who maintain their balance in the currency for which the remittance system is defined will be able to transfer funds.

3. Available transfer amounts for users of mobile operator services in various countries are provided by TransferTo in the currency of the destination country. Add the corresponding amounts on the

Destinations page in the following format: CCCNNN, where CCC is the currency code and NNN is the amount, e.g. IDR50000, INR500, CNY100 etc.

4. A specific tariff should be created for each currency assigned to a payment remittance system on the **Currencies** page. It will only be used for customers sharing the same currency. Create the following tariff:
 - **Name** – The tariff name depends on the currency; define the name in this format: TransferTo**CCC** (e.g. TransferToUSD) where CCC is the currency assigned to the corresponding payment remittance system.
 - **Currency** – Use the same currency as that used for the tariff name.
 - **Service** – Choose **Messaging Service** here.
5. Add rates to this tariff with the amounts your customers will be able to transfer:
 - **Rate Codes** – Choose the destination created in step 3 (e.g. IDR50000, INR500, CNY100, etc.) for the allowable top-up values.
 - **Unit Price Initial** – In this field, specify how much you want to charge the donor's account for this transfer.

Note that you can only set one fixed price for a 'currency'–'amount to transfer' pair, although TransferTo may charge differently for transferring funds to recipients in different countries or to recipients of different mobile operators within the same country.

For example, TransferTo will charge \$5.40 for transferring \$5.00 to a recipient of AT&T in the United States and \$6.20 for transferring \$5.00 to a recipient of BellSouth in Panama. Therefore, it is best to set up a price (the amount to be debited from the donor's account) which allows for the difference in transaction fees for various recipients.

Using the payment remittance system

Customers can transfer funds in one of two ways: via the account self-care web interface, or by a special application access number. The recipient is specified and identified by the phone number. The full list of amounts available for transfer is obtained from TransferTo, but only amounts present in the corresponding tariff that are within the customer's available balance are available for transfer.

14. Web interface administration

Multi-language support

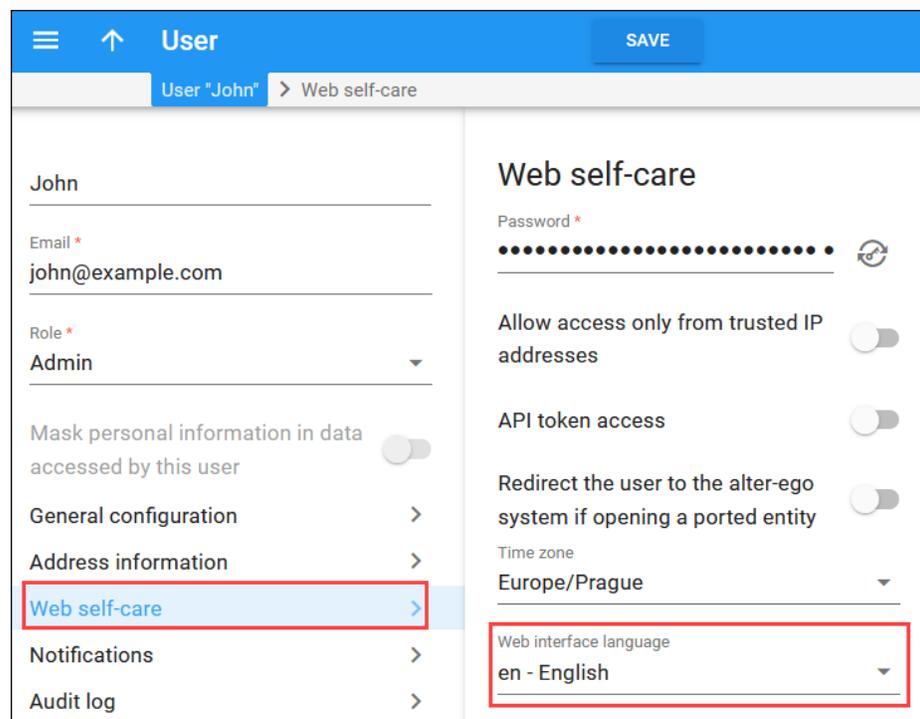
Though English is the main language used in PortaBilling®, other languages are also supported and administrators and customers can operate in them.

Depending on the type of user, these are the languages available.

User	Languages
Administrator	<ul style="list-style-type: none"> • English • Spanish
Customer	<ul style="list-style-type: none"> • Arabic • Chinese Simplified • Chinese Traditional • English • French • German • Hebrew • Italian • Norwegian • Polish • Portuguese Brazilian • Russian • Serbian • Spanish
Account	
Reseller	<ul style="list-style-type: none"> • Arabic • Chinese Simplified • Chinese Traditional • Czech • English • Estonian • French • German • Hebrew • Hungarian • Italian • Latvian • Lithuanian • Norwegian
Distributor	
Representative	
Vendor	

- | | |
|--|---|
| | <ul style="list-style-type: none"> • Polish • Portuguese • Portuguese Brazilian • Russian • Serbian • Slovenian • Spanish • Swedish |
|--|---|

Users can choose their language for their self-care interface. This does not affect other users' language settings. For example, John, an administrator, uses English on his web interface, while Peter uses Spanish as his web interface language.



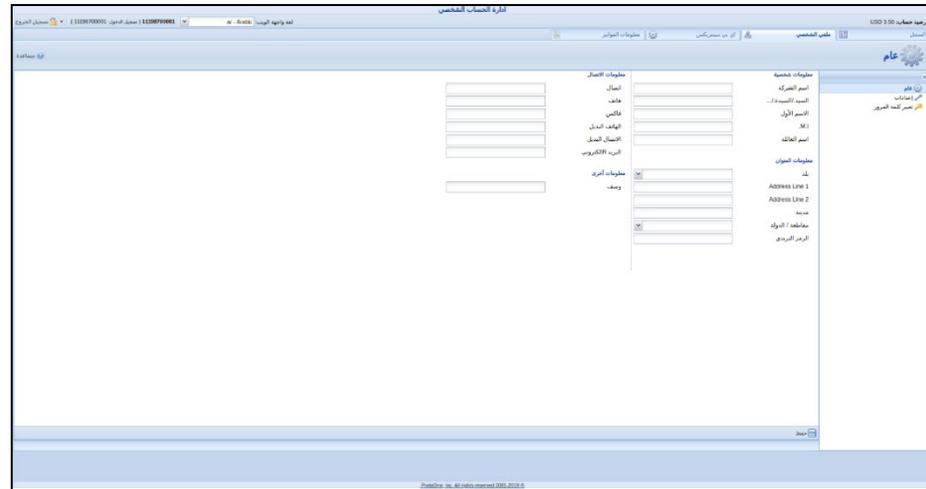
Translations are continually being updated and improved by the PortaOne customer community. We encourage PortaOne customers to make translations to other languages. ITSPs with the same background make translations targeted to their specific audiences since they use the same terminology. This kind of localization approach is very effective.

New languages can be added upon request. If you want to make and submit a translation, please contact the PortaOne support team to obtain the JS / XML template files.

Right-to-left layout in PortaBilling®

PortaBilling® supports right-to-left (RTL) layouts. The menu and the panels are mirrored and placed to the right and the text is aligned accordingly.

RTL layout is available for the customer and account self-care interfaces in Hebrew and Arabic.



This enhancement makes PortaBilling® user-friendly for customers who read right-to-left (e.g. Hebrew, Arabic) and improves their customer experience.

Access levels

Introduction to the ACL system

Different types of users have different responsibilities within the billing system. Some users may not be allowed to use or see certain portions of the system. To this end, PortaBilling® supports the concept of **Access Control Lists (ACL)**. ACLs allow the PortaBilling® administrator to decide, for example, that a particular sales representative can look at customers' data, but cannot create new customers.

ACLs allow you to control what users of your site can and cannot do. Without such restrictions, it is almost impossible to guarantee that users will see or change only the information that they are allowed to. There are default ACLs defined in the PortaBilling® system. You can use default ACLs or create new ones to fit your needs.

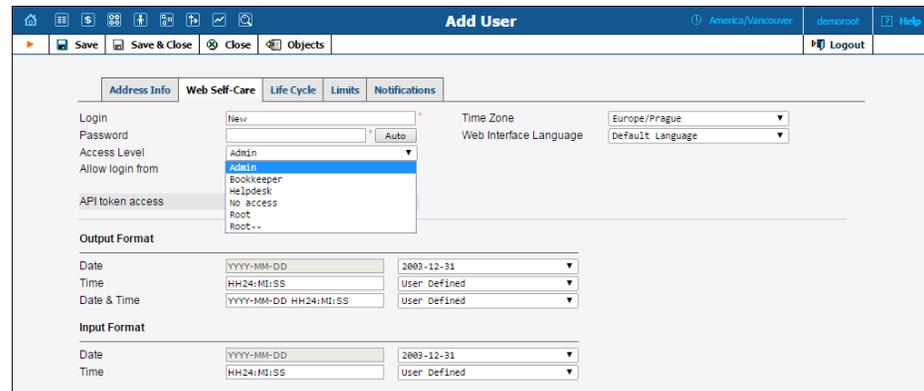
ACL types

An access level can be of the following types:

- Account (to be applied to your account)
- CC Staff (to be applied to your customer care support)
- Component (cannot be assigned to users; used only as a building block to construct other access levels)
- Customer (to be applied to retail customers or subcustomers)
- Distributor (to be applied to your distributor)
- Representative (to be applied to your representative)
- Reseller (to be applied to your resellers)
- User (access level for users of the admin interface)
- Vendor (to be applied to your vendors)

These access levels are composed of **permissions** and, optionally, other **components (as dependencies)**. Permission is a basic unit in the ACL system.

Newly created ACLs will be available in the **Access Level** select menu of the corresponding form when creating a new object or modifying an existing object's details. For instance, a User ACL will appear in the **Access Level** select menu of the **Add User** form (see below), a Customer ACL will be available when creating or editing a customer, and so on.

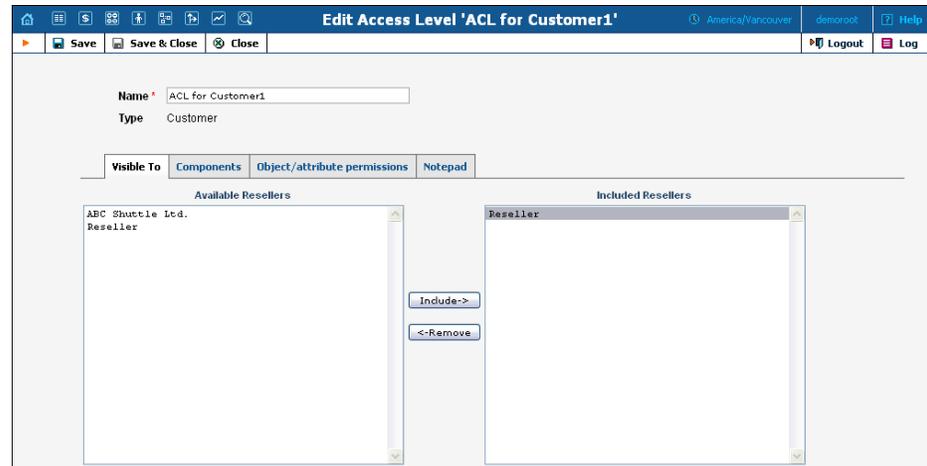


The screenshot shows the 'Add User' form with the following fields and values:

- Login:** New
- Password:** [Empty]
- Access Level:** Admin (dropdown menu is open showing: Admin, Bookkeeper, Helpdesk, No access, Root, Root--)
- Allow login from:** [Empty]
- API token access:** [Empty]
- Time Zone:** Europe/Prague
- Web Interface Language:** Default Language
- Output Format:**
 - Date: YYYY-MM-DD, 2003-12-31
 - Time: HH24:MI:SS, User Defined
 - Date & Time: YYYY-MM-DD HH24:MI:SS, User Defined
- Input Format:**
 - Date: YYYY-MM-DD, 2003-12-31
 - Time: HH24:MI:SS, User Defined

ACLs' Visibility under reseller

Normally you would not want reseller A to be able to use ACLs, which were designed for reseller B. To ensure that this never happens, ACLs are not visible to resellers by default. To allow a certain reseller to use the ACL, include this reseller in the ACL's **Visible To** tab:



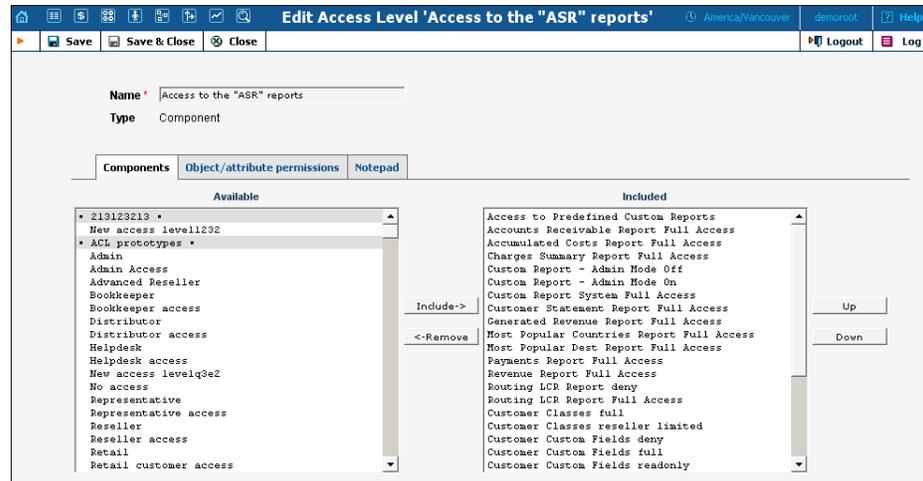
Visibility can be applied for customers', accounts', representatives' and CC staff's ACLs.

Components

As was mentioned before the system includes a set of default ACLs that consists of components. These are used as a building block for constructing other access levels. Components will be made up of zero or more permissions, and can include other components (as dependencies). If access level ACLX includes access levels ACL1, ACL2 and ACL3 (or, in other words, is *derived* from ACL1, ACL2 and ACL3), then ACLX will contain all the permissions defined in ACL1, ACL2 and ACL3 (along with all of the access levels they in turn are derived from).

What happens if there is a contradiction; for example, if ACL1 denies read access to **Accounts.password** and ACL2 grants it? In such a case, the first available definition will be used. Thus, in the example above, access will be denied according to ACL1, which is first in the list of included access levels. Keep in mind that the sequence of ACLs matching is held top-down. In addition to these, a component has several other aspects.

When editing a component, you will first see a page as on the following screenshot:



Including components as dependencies within other components gives the system its power. Here we see that the “Admin access” level is defined by over a dozen dependent components. Note that this component does not actually define permission itself, but rather relies on the implementation of its dependents.

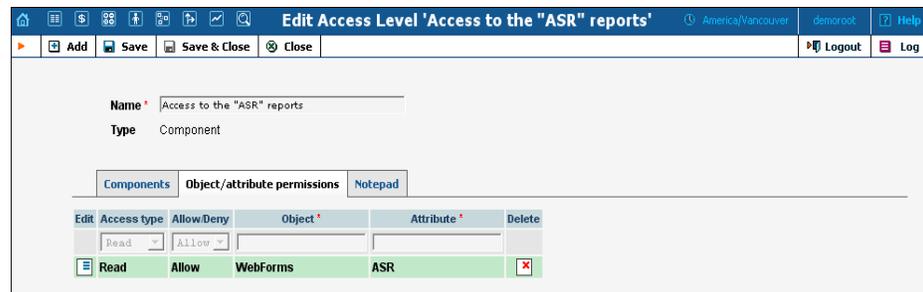
By deriving new components from the existing ones in the system, you can implement fine-grain access control and define User ACLs specific to your operational environment.

Permissions

Permission is the fundamental unit of exchange in the PortaBilling® security model. Permissions are composed of an **access type**, **Allow / Deny permission** (whether or not this is an allowed action), **the relevant object**, and **the relative attribute of the object**. To browse the permissions available for a certain page on the web interface, click the  **Objects** icon.

NOTE: The  **Objects** icon is available for Root users only.

Let’s take the example. An access level called “Access to ‘ASR’ reports” is provided within the PortaBilling® installation. It defines only one permission, which appears as in the following screenshot:



The “Access type” is set to “Read,” and the permission to “Allow.” This permission applies only to “WebForms” objects which have the attribute “ASR.”

There are four possible access types:

- **Read** – View the specified resource.
- **Update** – Modify the resource.
- **Insert** – Create new instances of the resource type.
- **Delete** – Remove instances of the resource from the system.

The **Allow/Deny** field defines whether this permission has been granted or withheld.

You should never have to provide fine-grain permission information yourself, as all possible permissions are already encapsulated in the components of your PortaBilling® installation. For this reason, we will not discuss the “Object” and “Attribute” fields further in this section. However, it may be useful to know that wildcards can be used in these fields. For example, to allow Read access to all web pages, an ACL could be defined with the following permissions:

- **Access type** – “Read”.
- **Allow/Deny** – “Allow”.
- **Object** – “WebForms”.
- **Attribute** – “*”.

As may be guessed, the “*” in the **Attribute** field means “all attributes.”

NOTE: Should you find it necessary to create new components, please contact PortaOne Support. We can provide you with a list of Object and Attribute resources, or otherwise assist you in maintaining ACLs.

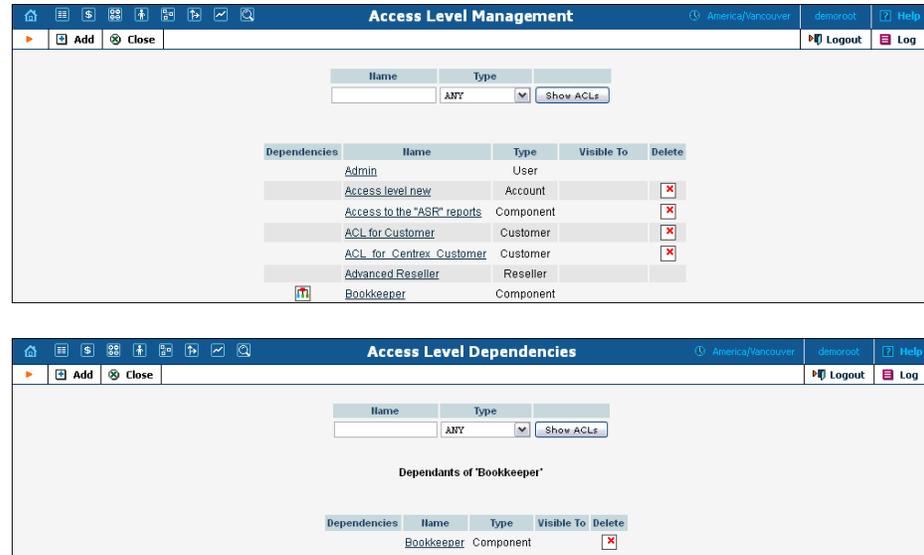
Access level management interface

In this discussion of the ACL system, we have proceeded by starting with the fundamentals and building up your skills from there. Now we will discuss the entry point for ACL management. On the PortaBilling® admin interface you will find a link to “Access Levels.” This link takes you to the Access Level Management main page.

This page is similar to many others in the PortaBilling® system, including a search interface at the top and a results listing at the bottom. You may search for ACLs using any combination of Name and Type.

In the results listing, you may also see the  **Dependencies** icon and the  **Delete** icon. ACLs can only be deleted when they are not in use.

If a component contains any included components, you will be able to click on the dependencies and see search results for all dependents. The following screenshot shows all dependents for “Accounts full access.” Note that some of the dependents have their own dependencies.



Add / edit a new user ACL

From time to time you will find that the PortaBilling® predefined user ACLs (Admin, Helpdesk, and so on) offer too few, or too many, restrictions for a particular class of user. In such a case, it is time to create a new user ACL.

The easiest method is to take an existing access level and create a new one modeled on it, and then modify it to fit your needs. You should examine the permissions granted to the model access level, and verify that you want to grant access to those resources.

Next, you can include other components to suit your needs. As a style recommendation, we suggest that you first create a component containing the dependent components you wish to utilize.

Finally, create a new user ACL which includes only this new component. Now you can assign this ACL to new users.

1. On the **Access Level Management** page, click **Add**.
2. On the **New Access Level** page, specify the following information:
 - **Name** – Type the custom name for this ACL.
 - **Type** – Select the entity this ACL is applicable for.
 - On the **Components** tab, define which components will be included in this ACL and in which order by using the **Include**, **Remove**, **Up** and **Down** buttons.

3. Click **Save**.
4. If required, on the **Object/Attribute Permissions** tab, define any additional permissions for this ACL.
5. Click **Save**.

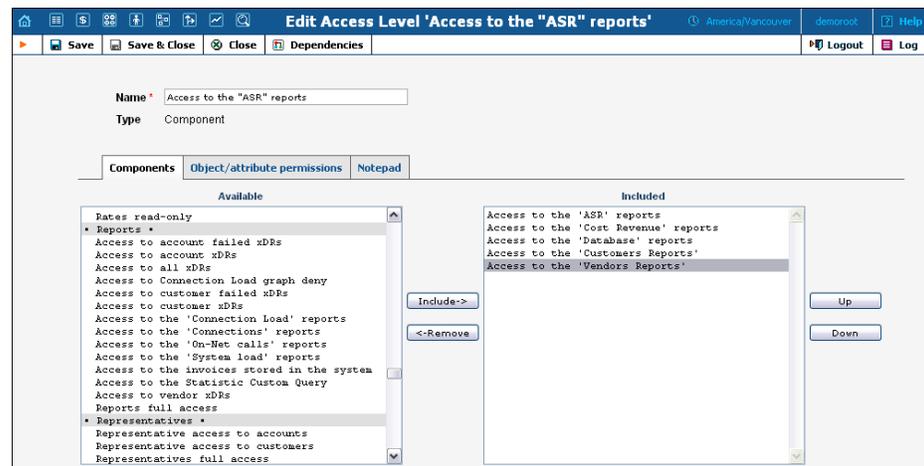
The PortaBilling® ACL management system contains style conventions which you would be well-advised to follow:

- The name of a component should be descriptive, based on the actions which it allows (for example, “Delete a node,” “Currencies read-only,” and “Access to Vendor Reports”).
- By convention, when defining a new user ACL (for example, “DemoUser”), we append “access” to the name of a component (“DemoUser access”) that includes dependent components.

Components tab

We have already talked about the necessary parameters for creating or editing components, but we have not yet discussed component inclusion in detail.

Each access level may have zero or more dependent components. These components are ordered, and likewise are applied in order until the first matching permission is reached. Keep in mind that the sequence of components matching is held top-down as shown on the screenshot:



In order to understand this better, we will use the previous example. Suppose a user is trying to view ASR reports. His access level must allow reading of “WebForms.ASR” (object “WebForms,” attribute “ASR”). For the sake of simplicity, we will say that his access level includes “A,” “B,” and “C,” where “B” allows this permission, but “C” explicitly denies it. In this case, the user’s ability to view these reports is based on the ordering of these components. If “B” appears before “C,” then it will work. In the opposite case, he will not have access.



This may sound complex, but in practice the user interface is quite simple. Two columns are shown on the **Components** tab of the edit page for each access level. On the left, you have a list of the available components, while on the right are the included components. Between these two columns you have the **Include->** and **<-Remove** buttons, which move selected items between the two lists. As for ordering, the **Up** and **Down** buttons on the far right-hand side of the page allow you to rearrange selected elements of the **Included** column.

You should now have the skills necessary to implement the PortaBilling® security model and customize it to suit your business environment.

Role-based access control

Along with the ACL security system, administrators can use the role-based security system to control user access to all of the resources in PortaBilling®. Access control lies in configuring a role and assigning this role to a user. This ensures that the user can access only those resources they are authorized to see or use.

The role-based security system is accessible from the new web GUI and exists in parallel with the ACL one. Administrators can use both security systems to control user access to application resources.

Roles

Default roles are supplied with PortaBilling® – or administrators can create new roles to fit your company needs.

A role can be one of the following types:

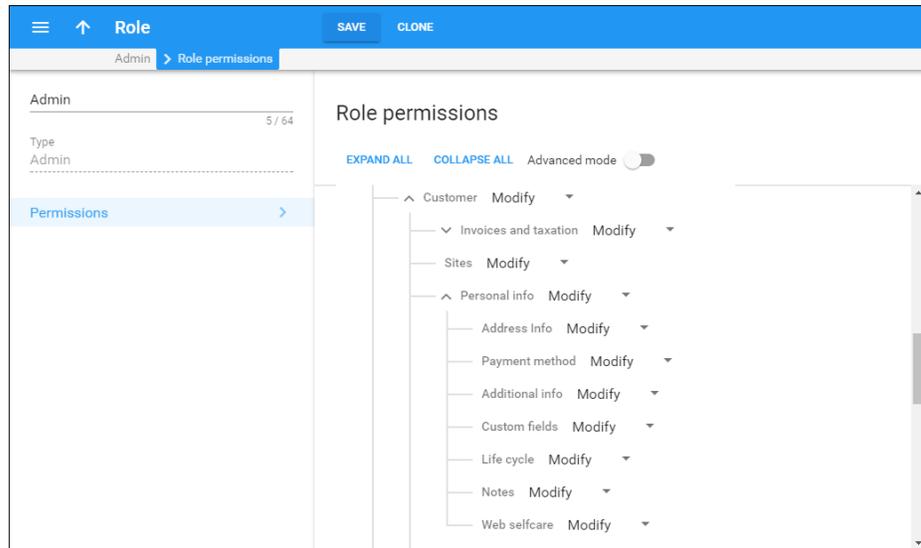
- Account – to be assigned to accounts
- CC Staff – to be assigned to customer care support employees
- Customer – to be assigned to retail / reseller customers
- Distributor – to be assigned to distributors
- Representative – to be assigned to representatives
- Reseller – to be assigned to resellers
- Admin – to be assigned to users of the admin interface
- Vendor – to be assigned to your vendors

For now, the Admin, Account and Customer role types are available. Support for other role types will be implemented in subsequent releases.

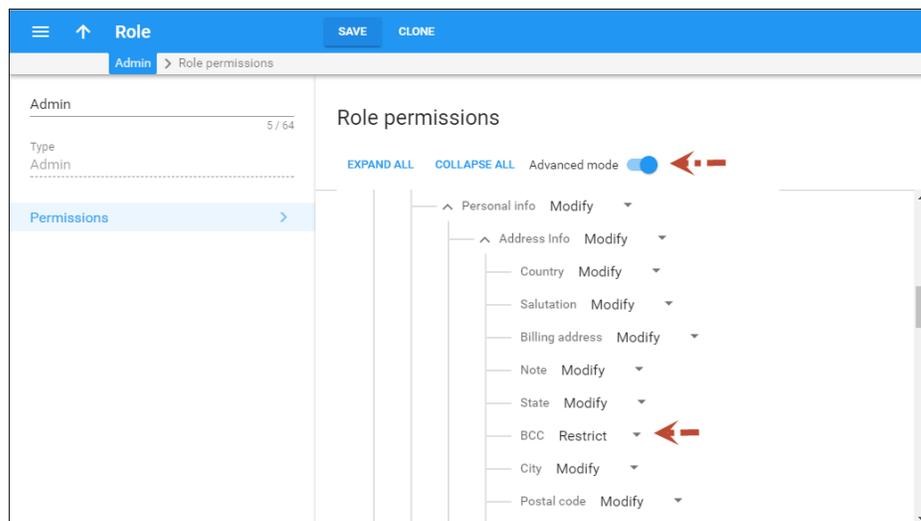
Roles are presented on the web interface as a resource tree wherein root nodes reflect entities in PortaBilling® (i.e. customer, account, product, etc.). Second-level nodes reflect entity parameter panels. For example, for

a customer entity, it can be the **Personal information** panel, **Invoices and taxation** panel, etc.

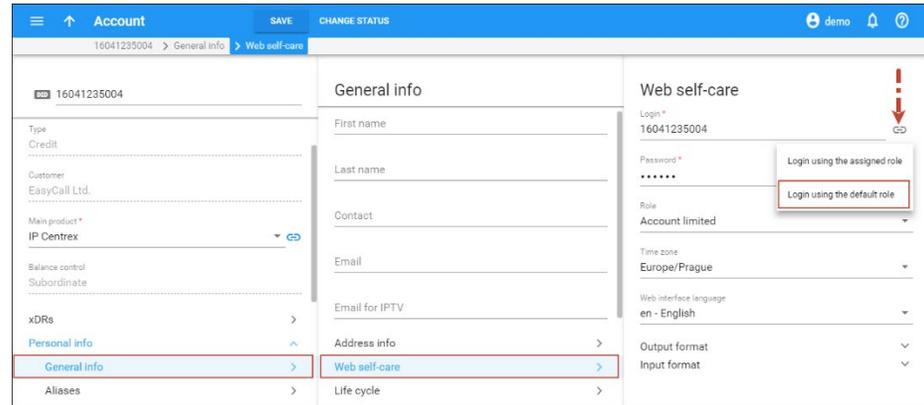
For each node within the tree, the administrator assigns permissions to define whether an entity or its parameters are available for the user and which actions the user can perform on them. The role’s resource tree has a hierarchical structure, that is, lower-level nodes inherit permissions that have been assigned to higher-level nodes.



If the administrator needs to hide a certain item on an entity parameter panel (e.g. the **BCC** field on the **Address info** panel), they can switch to the Advanced mode. In this mode, the role’s resource tree displays a list of items for each entity parameter panel. The administrator then sets the required permission for the corresponding item.



The administrator uses roles to control access to the self-care page for customers / account owners. The administrator can log into a customer self-care page on behalf of the customer, for example, to help with the IP Centrex configuration. If the customer has a custom role assigned to them, the administrator can choose whether to log in with either the custom or the default role.



Permissions

The administrator can assign one of the following permissions in the role's resource tree:

- Restrict – This means that users cannot access the specified resource.
- Read – This permits users to view the specified resource.
- Modify – This permits users to view, update, create and delete the specified resource.

When a user attempts to perform a specific action with a resource (for example, update customer information), PortaBilling® checks whether the user has permission for this action. If permission is granted for this action, the user may proceed. Otherwise, the action is not permitted.

Two systems operating in parallel

The role-based and ACL security systems operate in parallel in the following way.

ACLs are accessible in both systems. In the role-based security system, administrators can assign users both roles and ACLs. However, administrators cannot see or modify ACL details.

Roles are only accessible in the role-based security system. In the ACL security system administrators can neither see / modify role details nor assign roles to users.

What happens if there is an inconsistency? For example, one administrator assigns ACL1 (or a role) to a user in the role-based security system and another administrator assigns ACL2 to this same user in the ACL system? In this case, whatever the last changes are, are the changes that will go into effect.

When the user logs in to the new web GUI the system checks their role permissions. When the user accesses the old web interface, his ACL permissions are checked before the user is granted access to the system.

The same logic is in place when operating with the system via the API. The components that were already moved to the new web GUI are covered by role permissions. This means that when an API user sends a request to retrieve a customer list or edit a product, the system checks the user's role permissions before executing the request. If the user wishes to modify the component that is still in the old web interface, the system checks the user's ACL permissions.

During a system update, users with system default ACLs are automatically assigned roles to access the new web GUI. Corresponding roles are created for users with custom ACLs, though an administrator must assign them to users. The corresponding custom roles for customer / account self-care pages are created and assigned automatically.

Limits management for administrative users

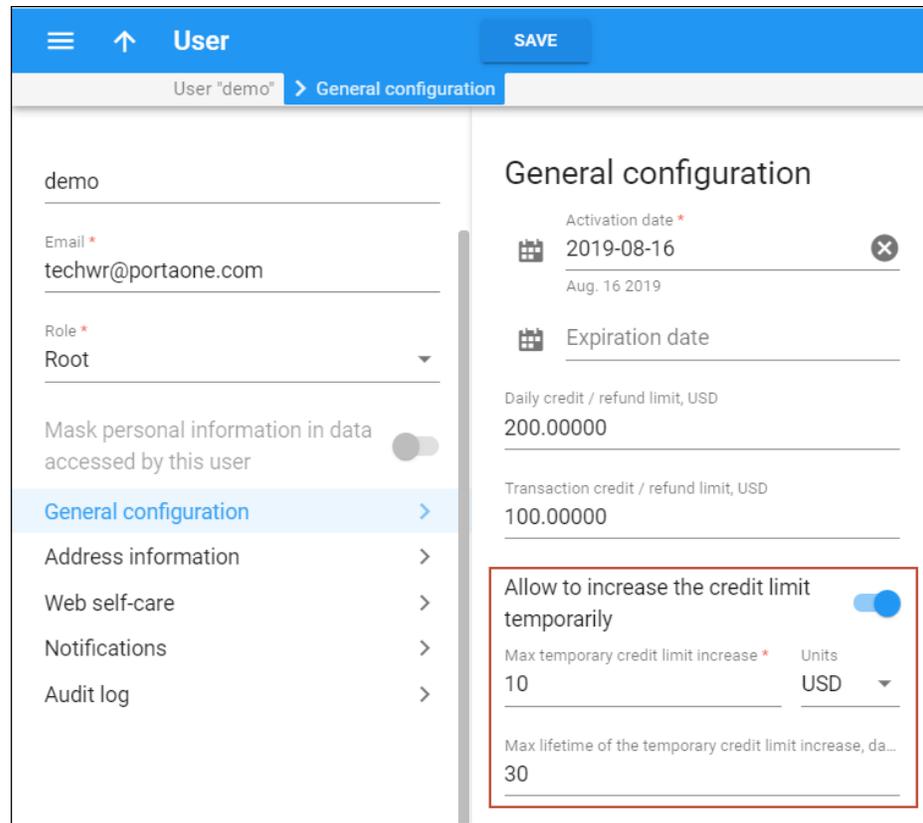
Temporary credit limits

You can empower employees to define a temporary credit limit for a customer depending on the confidence level that you share with your employees. You can:

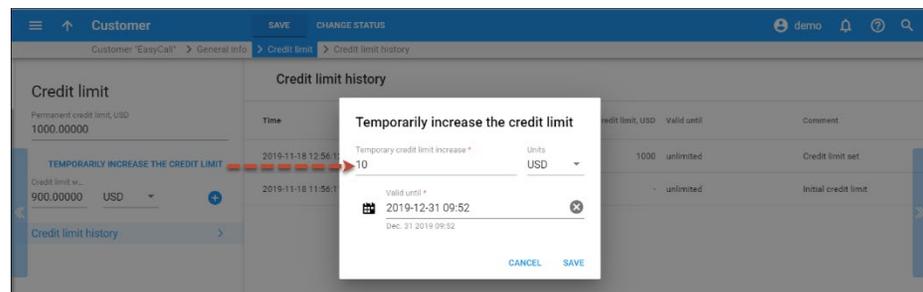
- Adjust a temporary credit limit for each PortaBilling® user depending on the trust level you share with them.
- Assign a temporary credit limit amount.
- Set an amount either as an absolute value or as a percentage of the original credit limit.
- Assign a period for temporary credit limit usage.

To do this, enable the **Allow to increase the credit limit temporarily Amount** option on the General configuration panel of the administrative user. Specify:

- The **maximum amount** of the temporary credit limit and
- The **duration** for this temporary credit usage.



Your staff can increase the credit limit for a particular customer, distributor and reseller on the **Credit limit** panel. The amount and duration of this temporary credit limit cannot be higher than allowed for this user.



Consider the following example: A business owner sets a credit limit increase for a Helpdesk operator named Roger as follows:

- Temporary credit limit amount – 10% of the original credit limit;
- Duration – 30 days.

For the customer relations manager (CRM), Andrew, he defines different values:

- Temporary credit limit amount – 20% of the original credit limit;
- Duration – 60 days.

Customer John Doe has a \$200 original credit limit. He has reached the limit and he has made a payment, however, the money has not yet transferred to his PortaBilling® account. John urgently needs to make some calls so he attempts to obtain a temporary credit limit increase. He calls the ITSP Helpdesk and asks to recharge his account by \$20 for the next 7 days. The helpdesk operator satisfies this request since both parameters (amount and duration) do not exceed the maximum values allowed for the helpdesk staff.

The customer EasyCall Ltd. has a \$1000 original credit limit. The company also faces a situation when their original credit limit is reached. The representative of the company makes a call to the ITSP Helpdesk and requests a credit limit increase of \$200 for 40 days. The helpdesk operator does not have permission to raise the limit (neither the amount nor the duration are allowable), so he raises the issue to the CRM. The CRM knows that EasyCall Ltd. is a reliable customer that can be trusted, so he increases the credit limit and the duration period. The customer EasyCall Ltd. is satisfied and continues to enjoy the services.

Consequently, it is now possible to define different temporary credit values for CRMs, helpdesk operators and other employees, individually. A business owner can easily define and change these values on the PortaBilling® web interface; default values are set on the Configuration server.

Daily and transaction limits on credit / refund for PortaBilling® users

To prevent potential abuse by employees and avoid money leaks, you can define daily and transaction limits on credits / refunds for PortaBilling® users.

On the **General configuration** panel, you configure the following:

- **Daily credit / refund limit** – This defines the amount the user cannot exceed for generating either a credit or a refund for customers during a 24-hour period.
- **Transaction credit / refund limit** – This defines the amount the user cannot exceed per transaction.
- **Max temporary credit limit increase** – This defines the maximum amount by which the user can increase the temporary credit limit (either as an absolute value or as a percentage of the permanent credit limit), and the credit usage's valid time span.

The screenshot shows the 'User' configuration page for 'User "Pete"'. The page is divided into two main sections: a left sidebar with navigation options and a main content area for 'General configuration'. The sidebar includes 'Pete', 'Email *', 'Role *', 'Mask personal information in data accessed by this user', 'General configuration', 'Address information', and 'Web self-care'. The main content area includes 'Activation date *', 'Expiration date', 'Daily credit / refund limit, USD', 'Transaction credit / refund limit, USD', and 'Allow to increase the credit limit temporarily'. The 'Daily credit / refund limit, USD' field is highlighted with a red box and contains the value '200.00000'. The 'Transaction credit / refund limit, USD' field contains the value '10.00000'. A 'SAVE' button is located in the top right corner.

Limit values change after each transaction that impacts the limit: manual credit, promotional credit, refund and e-commerce refund. The limits are reset at midnight (in the time zone of the PortaBilling® user's environment).

Let's consider the following example. The business owner defines a daily limit of \$200 and a transaction limit of \$10 to helpdesk operator Pete.

- Pete attempts to apply a \$15 credit to a customer but the operation fails.
- Pete applies a \$3 credit to a customer and the operation is successful, so Pete's daily credit / refund usage is now \$3.
- Pete applies a \$2 refund to a customer and the operation is successful, so Pete's daily credit / refund usage is now \$5.
- Pete applies a \$20 manual payment to a customer and the operation is successful, so Pete's daily credit / refund usage does not change (it is still \$5).
- Pete applies a \$20 credit card charge to a customer and the operation is successful, so Pete's daily credit / refund usage again does not change (it is still \$5).
- Pete subsequently applies a \$10 credit to 19 different customers and all of the operations are successful, so Pete's daily credit / refund usage is now \$195.
- Pete attempts to apply a \$10 credit to a customer but the operation fails since Pete's daily limit is exceeded.
- After midnight, Pete's daily credit / refund usage is again reset to its defined values and he can retry the transactions.

Thus, you protect your business via money allocation control in the credit arena.

Quick forms

A quick form is used to simplify and quicken the process of creating new customers and accounts and avoid mistakes when filling in parameters. The administrator fills in most of the important details (e.g. customer class, currency or product for an account), so that later the person entering data only has to input a few remaining ones (e.g. customer name) to create a customer and an account (or multiple accounts) under it.

Custom fields

Custom fields are used for storing additional attributes that supplement standard PortaBilling® information. This can be information about a customer's bonus, driving license ID, or a code name that maps an Internet customer to a certain vendor, etc. This can also be very useful for troubleshooting. Create your own custom fields and give them whatever name you like, set a field type and then control their visibility, etc.

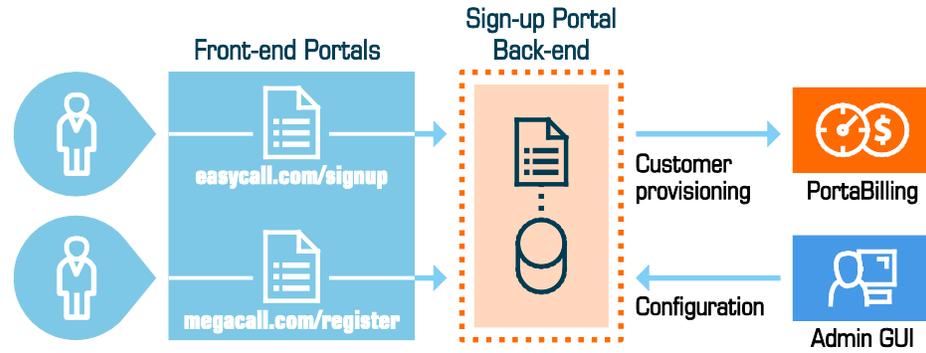
Custom fields are treated like any other field; you can organize them via customer classes and fine-tune their visibility for end users. Then desired custom fields can be set on the **Customer / Account Info** page and used for search queries.

Custom fields are added to **Invoice** templates under the **Header** section and it is possible, as well, to add custom fields in Layout Designer (right-click→Header).

This functionality allows you to manage extra user information that in its turn facilitates customer support and helps with troubleshooting.

Online web signup

PortaBilling® allows you to automate the process of subscribing new customers to the service, so that customers can do everything necessary online without requiring any assistance from your administrator. A template-based API for online signup allows you to customize page layout and put a web signup front-end on multiple websites (web portals, web pages of your resellers or affiliates, etc.). The following diagram illustrates the online web signup process:



A fully functional web signup portal is supplied with PortaBilling®. You can use it out of the box, or take it as the basis for further customizations.

Web signup allows you to:

- Collect customer name and address information.
- Allow customers to select from the available phone numbers.
- Choose a desired product bundle.
- Collect a customer's credit card information.
- Authorize credit card payments made via payment systems that don't require external authorization, fully provision the account information in PortaBilling®, and provide service-related information (e.g. auto-configuration file for a softphone).

This is a three-level application which allows you to deploy multiple online web signup sites without compromising security, and to easily customize each site to fit a specific business case.

- Front-end – These are the web pages customers will access using their web browser. It is a set of HTML pages (with Javascript and CSS functionality), which can be fully customized in terms of page layout, colors, etc.
- Middleware – This is the code located on the PortaBilling® server or another secure server. It retrieves the required data from PortaBilling® using XML API and communicates with the front-end using AJAX. This serves as a protection layer against any malicious web requests.
- Back-end – Located on the PortaBilling® server, it performs credit card authorization and actual data modification in the database.

Personal data protection. GDPR compliance

GDPR (General Data Protection Regulation) is the EU regulation concerning personal data protection for EU residents. GDPR also

regulates the transfer of personal data for processing by third parties outside the EU.

Personal data is information by which a person can be identified. Personal data includes name, address, phone number, MAC address of an IP phone, etc. Find the full list of personal data (defined by the administrator or provided by customers themselves) that may be stored in PortaBilling® in [Appendix G](#).

To help you comply with this regulation the following enhancements are introduced in PortaBilling®:

- Anonymizing personal data on the PortaBilling® web interface and in API responses;
- Two-version PDF invoices and custom reports: one that contains full information and another that contains only anonymized personal data;
- Adjusted email / SMS notifications received by administrators to anonymize personal data;
- Recording user access to personal data in logs for troubleshooting and auditing purposes; and
- The “Right to be forgotten” functionality – the ability to completely remove customer personal data.

Personal data anonymization and access control

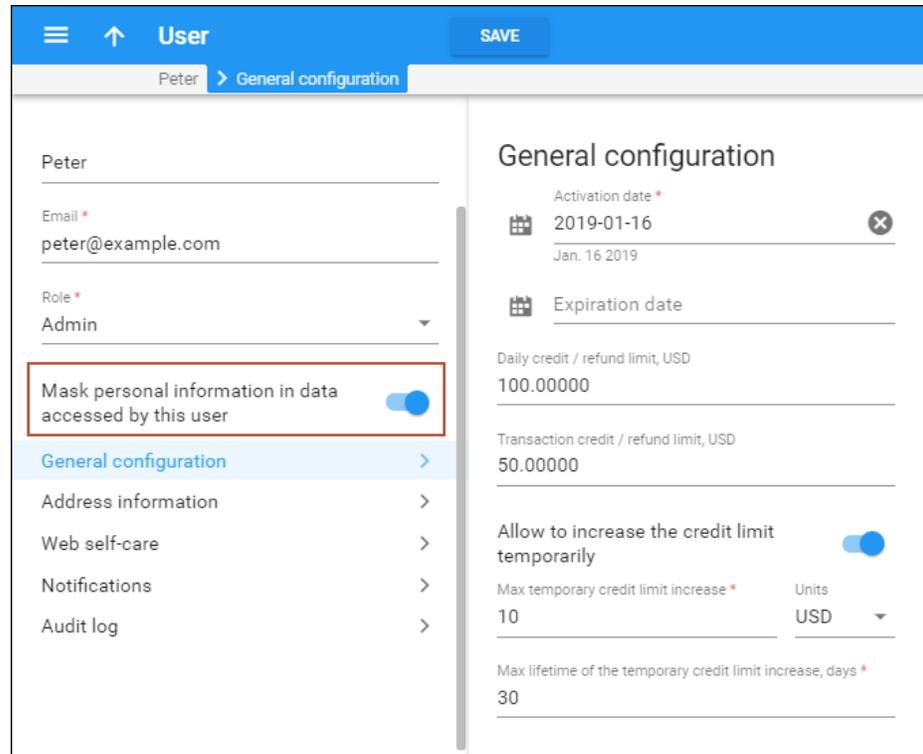
You can divide your administrative staff members into those who can see full customer details and those who are not allowed to process personal data.

Consider the following example:

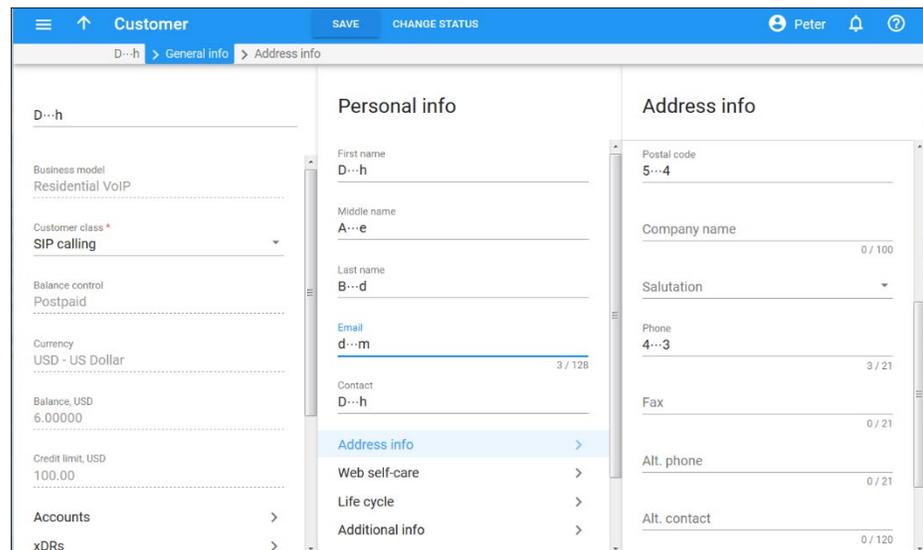
Let’s say you provide services in France and your staff is in Paris, plus you have a remote administrator, Peter, in Montreal.

Peter must not process personal data without user consent and must not have access to it. Since you find it difficult to meet the GDPR requirements regarding personal data transfer and processing outside the EU, you decide to restrict Peter’s access to personal data.

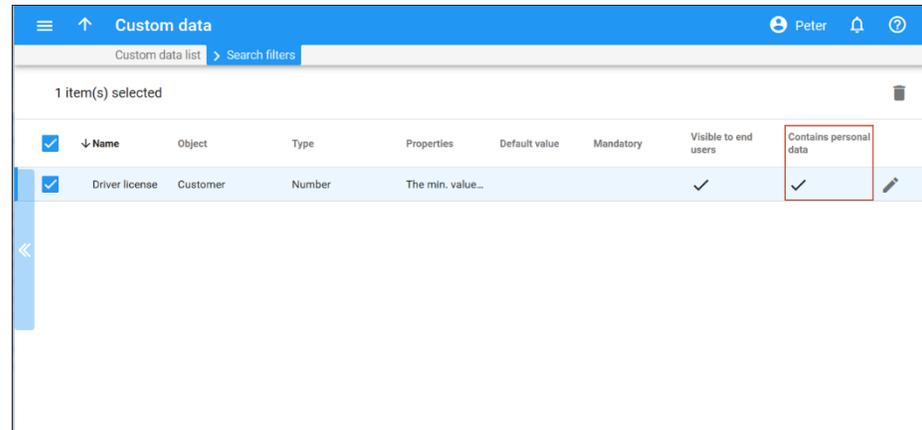
To do this, you enable the **Mask personal information in data accessed by this user** option for his user in PortaBilling®.



This way, when Peter opens a customer’s page, he sees that personal data is anonymized. When Peter receives an email or SMS notification that an invoice for a customer is re-generated, the customer name in the notification is also anonymized.



If you store additional information about your customers in custom fields and it could be considered personal (e.g. driver’s license or insurance ID), you need to anonymize it, too. To make this happen, mark this field as containing personal data in PortaBilling®.



Peter can browse and modify customer details (e.g. assign add-ons to a customer account) and create new customers and accounts. However, upon customer creation or modification, he sees that the customer's personal data is anonymized

All personal data, such as phone numbers, IP addresses and ports of call participants is anonymized in SIP and BE logs to prevent Peter's access to it when he troubleshoots issues for a customer.

Peter's activity in PortaBilling® has the following limitations:

1. Access to customer and account self-care portals is forbidden.
2. Attachments with personal data that cannot be anonymized, such as call recordings and music on hold files are not available.
3. All other attachments such as account generation reports, invoices in invoice notifications, files with DID, etc. are available since personal data in them is anonymized.
4. Access to SIP and BE logs is forbidden since the logs contain personal data. (Personal data anonymization will be added in future releases to grant access to logs.)
5. Previously generated custom reports are not available for download. Reports yet to be generated are available since personal data in them is anonymized.
6. Access to the old web interface is forbidden so that access to unanonymized personal data is prevented.

The same logic applies when your staff manages personal data in PortaBilling® via the API from an external self-care portal.

By default, all personal data is anonymized using the same pattern – only the first and last symbols are shown. In future releases we plan to introduce the capability to customize personal data masking.

Access control to personal data for CC staff members

Resellers can restrict access to their customers' personal data for CC staff members via the API if they must not process it according to the GDPR. To make this happen, your ABC reseller enables the **Mask personal information in data accessed by this user** option, for example, for their CC staff member Bob Smith.

The screenshot shows the 'Edit CC Staff' interface for a user named 'bobsmith'. The 'Access Level' is set to 'CC Staff'. The 'Mask personal information in data accessed by this user' option is set to 'Yes', which is highlighted with a red box. Other settings include 'Time Zone' set to 'Europe/Prague' and 'Web Interface Language' set to 'Default Language'. The interface also shows sections for 'Output Format' and 'Input Format' with various date and time format options.

When Bob logs in to his self-care portal and retrieves customer information via the API, all personal data (e.g. name, contact details, etc.) is masked.

Two-version PDF invoice files

To protect personal data, PortaBilling® produces two versions of customer invoices:

- Invoices with full data are generated for customers' use;
- Invoices with anonymized personal data are available for administrators who have restricted access to personal data.

Previously generated invoices are not available for download for these administrators either; however, they can adjust and re-generate such invoices, if necessary.

Logging user access to personal information

GDPR establishes requirements to record access to personal data by any user and consequently, to notify a supervisory authority should a personal data breach occur. Thus, service providers must track who processed personal data and when, and they must be able to prove it.

To meet these requirements, PortaBilling® records every user action that deals with personal information. So whenever an administrator with full access to personal data modifies a customer, the action is recorded in a log. Similarly, when a helpdesk operator with restricted access opens a customer panel, PortaBilling® logs the event.

The screenshot shows the 'Audit log' interface. On the left, there is an 'Audit log search' panel with various filters: 'From time' (2019-01-15 12:30), 'To time' (2019-01-18 12:30), 'Entity type' (Customers), 'Action type' (Read), and 'Property name' (Any). The main area displays a table of results with columns: Timestamp, User #, Entity type, Entity name, Action type, and Description. The table contains 12 rows of log entries, all showing 'Read' actions on 'Customers' entities.

Timestamp	User #	Entity type	Entity name	Action type	Description
2019-01-15 14:40:32	I-r 192...-42	Customers	I-r	Read	Personal data was read using the Cust...
2019-01-15 14:40:18	I-r 192...-42	Customers	I-r	Read	Personal data was read using the Customer_get_customer_info method
2019-01-15 14:39:42	p-t 192...-88	Customers	C-9	Read	Personal data was read using the Cust...
2019-01-15 14:39:42	p-t 192...-88	Customers	B-0	Read	Personal data was read using the Cust...
2019-01-15 14:39:42	p-t 192...-88	Customers	C-1	Read	Personal data was read using the Cust...
2019-01-15 14:39:42	p-t 192...-88	Customers	B-1	Read	Personal data was read using the Cust...
2019-01-15 14:39:42	p-t 192...-88	Customers	C-8	Read	Personal data was read using the Cust...
2019-01-15 14:39:42	p-t 192...-88	Customers	C-4	Read	Personal data was read using the Cust...
2019-01-15 14:39:42	p-t 192...-88	Customers	C-t	Read	Personal data was read using the Cust...
2019-01-15 14:39:42	p-t 192...-88	Customers	U-s	Read	Personal data was read using the Cust...
2019-01-15 14:39:42	p-t 192...-88	Customers	F-1	Read	Personal data was read using the Cust...
2019-01-15 14:39:32	p-t 192...-88	Customers	I-s	Read	Personal data was read using the Cust...

This provides the ability to demonstrate the proper handling of personal data from creation to deletion. You can also provide proof of who exactly accessed personal data and a time stamp.

Complete removal of customer personal data (“Right to be forgotten”)

To comply with the GDPR “right to be forgotten” requirements, an administrator can completely remove customers’ personal data from PortaBilling®.

The following personal data is removed from the PortaBilling® database:

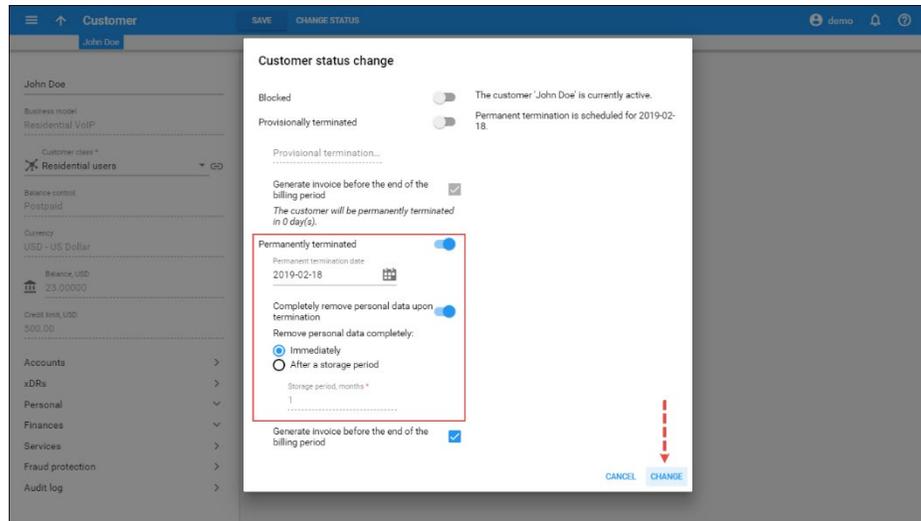
- All information about a customer and their accounts which is considered personal (e.g. name, billing and contact information, MAC address of an IP phone, etc.);
- all call recordings.

Customer and account xDRs remain in PortaBilling®, however, account IDs, CLIs, CLDs and IPs (e.g. subscriber_ip or originating_ip) are removed from them.

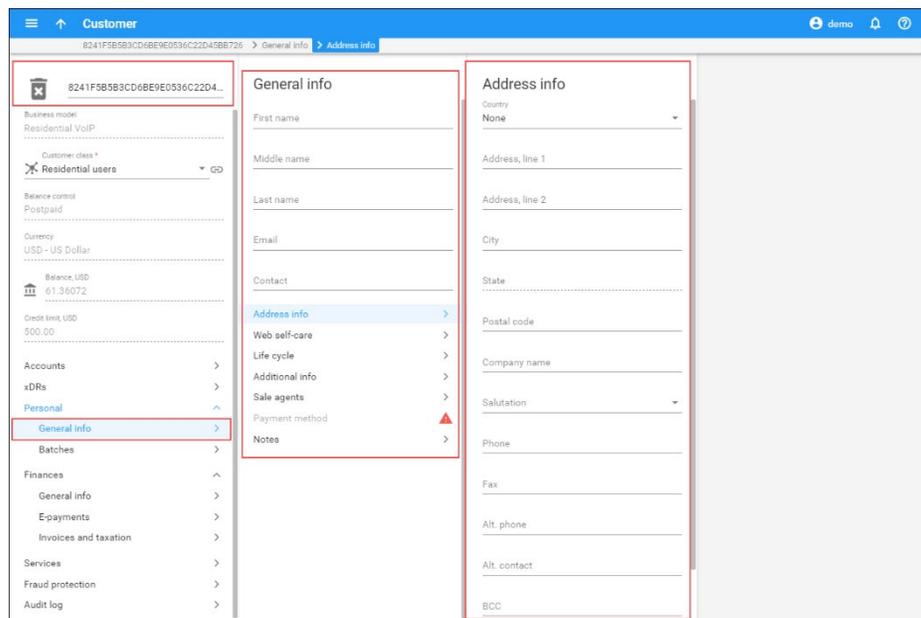
There are two options for removing a customer’s personal data:

- immediately upon termination; or
- after the specified storage period ends. (e.g. you store customer’s personal data until the limitation period expires).

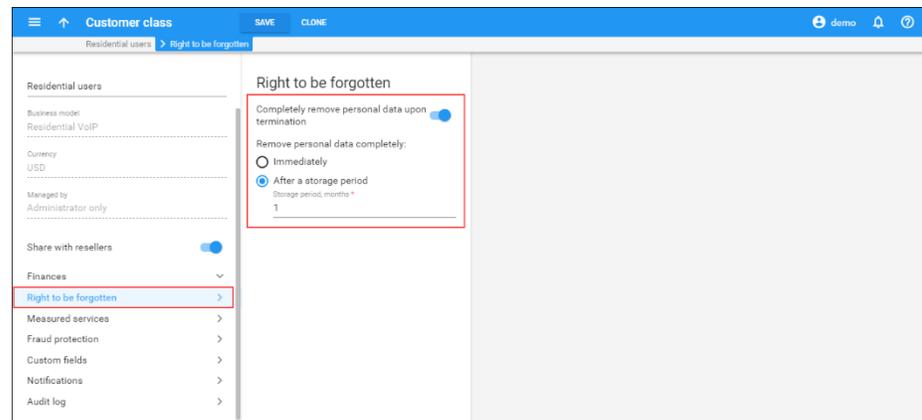
Let’s say customer John Doe asks the administrator to close his contract and remove all his personal data from the system. The administrator schedules the termination of John’s record and arranges for his personal data to be removed immediately upon termination.



So once John’s record is terminated, PortaBilling® deletes his personal data.



An administrator can preconfigure personal data removal via the customer class. Thus, all the customers of this customer class share the same configurations for data removal.



For example, an administrator can assign a “European Customers” customer class to all customers located in Europe and specify that these customers’ personal data be stored for one month.

White labeling

Using the brandpane feature you can customize your system on a per-environment basis to display your company name and logo on every web page.

Every realm within the environment (e.g. admin, customer, vendor, cc_staff, etc.) can be modified by using a separate template. Brandpane templates are stored in the database and are periodically cached to `~porta-admin/apache/brandpane`.

To use the brandpane feature you will need to create an appropriate template file hierarchy. Do NOT use `~porta-admin/apache/brandpane` for that as it is periodically updated by the caching procedure. Instead, use any temporary directory (e.g. `/home/demoroot/tmp/bp`).

Brandpane file hierarchy

- First-level sub-directories represent environments. The name of a sub-directory must be ‘default’ (for an unspecified environment) or ‘env_#’ (where # is i_env number). Sub-directory ‘default’ is required.
- Second-level sub-directories represent realms, named as in configuration (admin, accounts, cc_staff, customer, vendor, representative, etc.). Sub-directory ‘admin’ is required.
- Every sub-directory must contain at least one file named ‘brandpane.tpl,’ which is an ordinary HTML file. Additional files (images and CSS) are allowed.

```
DIR
default
  admin
    brandpane.tmp1
    my.css
    logo.jpg
  accounts
    brandpane.tmp1
    brandpane.css
  cc_staff
  ...
  customer
  vendor
env_1
  admin
    brandpane.tmp1
    another-logo.jpg
  customer
env_2
...
env_32
  admin
  ...
```

Brandpane template selection order

Brandpane is searched and applied in the following order:

- env_\${I_ENV}/\${REALM}
- default/\${REALM}
- env_\${I_ENV}/admin
- default/admin

`${I_ENV}` and `/${REALM}` are the current environment number and name of the realm. Steps 3 and 4 are not executed if the current realm is 'admin.'

Create new brandpane templates

Follow the steps below to create a new brandpane template:

1. The template's file hierarchy should be created in a temporary directory (e.g. in `/home/demoroot/tmp/bp`). Log in to the server with the admin instance assigned to it and type:

```
cd /home/demoroot/tmp/
mkdir bp
```

2. Create a proper brandpane file hierarchy and HTML components according to your requirements. Edit files `brandpane.tmp1` (a sample file is shown below).

```
<table width=100% border="0" align="left"
cellpadding="0" cellspacing="0" bgcolor="white">
  <tr>
    <td width="100%">
      
    </td>
  </tr>
</table>
```

Note that images must include a full path, relative to the Document Root (e.g. `"/brandpane/default/admin/mylogo.jpg"`).

3. Import the brandpane into the database:

```
impexp_templates.pl -i -brandpane --  
path=/home/demoroot/tmp/bp
```

4. Activate the brandpane feature on the Configuration server web interface (the **BrandPane_Enabled** option).

Modify existent brandpane templates

To change an already existent brandpane template, export it from the database to a temporary directory, make your desired modifications and then import it back to the database:

1. Export brandpane templates into temporary directory:

```
impexp_templates.pl -e -brandpane --  
path=/home/demoroot/tmp/bp
```

2. Edit files, add / remove images, styles, etc.:

```
nano /home/demoroot/tmp/bp/default/admin/brandpane.tpl
```

3. Import brandpane templates back into the database:

```
impexp_templates.pl -i -brandpane --  
path=/home/demoroot/tmp/bp
```

4. Wait until Apache updates the cached brandpane templates.

5. Clear the temporary directory:

```
rm -rf /home/demoroot/tmp/bp
```

Remove existent brandpane templates

To remove already existent brandpane templates, find the appropriate template in the database and remove it:

1. Find the brandpane template number in the database:

```
impexp_templates.pl -l -db
```

A command output example is shown below. Note that the template's name is also its realm.

```
...  
Row 121  
Type: Invoice  
Authorship: Custom  
Name: 228  
Description:  
Revision: 1.5  
  
Row 122  
Type: BrandPane  
Authorship: Custom  
Environment: 13  
Name: vendor  
Description:  
Revision: 1.9
```

2. Remove the template with the appropriate number:

```
sudo -u -www impexp_templates.pl -d -t=122
```

3. Remove the template from the file system on every admin instance of the web cluster:

```
rm -r /home/porta-admin/apache/brandpane/$I_ENV/$REALM
```

`$I_ENV` and `$REALM` are the current environment number and name of the realm.

Brandpane for self-care interfaces

The brandpane functionality for self-care interfaces works similarly to other realms. It is configurable and can be turned on / off on the Configuration server via the **ShowBrandPane** option. The only difference between them is that the brandpane template file for self-care interfaces must be in the HTML format. An example of the brandpane.html file is shown below:

```
<!DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 4.01  
Transitional//EN"**  
"http://www.w3.org/TR/html4/loose.dtd">  
<html>  
<head>  
  <meta http-equiv="Content-Type" content="text/html;  
  charset=utf-8">**  
  <meta http-equiv="X-UA-Compatible" content="IE=8">**  
<title>Account Self-Care</title>  
</head>  
<body style="background-color:Tomato">  
  
</body>  
</html>
```

The look and feel of the self-care interfaces' brandpane can be adjusted by modifying the CSS entries in the brandpane.css file (which should be linked to the brandpane.html file). For example, the following code changes the default brand pane height:

```
.brandPane{
```

```
    height: 200px;
  }
```

If you need to customize the self-care interfaces' brandpane in different ways for different customers or accounts, you can do this via the following variables: *i_customer* and *i_account*. An example below shows how this can be done:

```
<!DOCTYPE HTML PUBLIC "-//W3C//DTD HTML 4.01 Transitional//EN"
"http://www.w3.org/TR/html4/loose.dtd">
<html>
<head>
  <meta http-equiv="Content-Type" content="text/html; charset=utf-8">
  <meta http-equiv="X-UA-Compatible" content="IE=8">
  <title>Account Self-Care</title>
</head>
<body style="background-color:Tomato">
  <table>
    <tr>
      <td>
        <td>
          <h1> Hello customer # <%=i_customer%> </h1>
        </td>
        <td>
          % my $result = $i_customer % ( int(rand(9)) + 1);
          <% $result %>
          % if ( $result % 2 == 0 ) {
            
        </tr>
      </table>

</body>
</html>

<%args>
  $i_customer => undef
  $i_account => undef
</%args>
```

Thus, different logos will be displayed, depending on which customer has logged in to their self-care interface.

Adding a reseller's logo to the reseller's customers' self-care interfaces

1. Activate the brandpane feature as described in the topic above.
2. Configure an additional virtual web server which will run on the reseller's domain name (e.g. www.supercall.com).
3. Include conditional statements in the brandpane component. These will analyze the domain name the customer is trying to access, and will then produce the company name and logo for your domain (www.sipcall.com), while for the reseller's (www.supercall.com) the reseller's name and logo will be produced. See the following example:

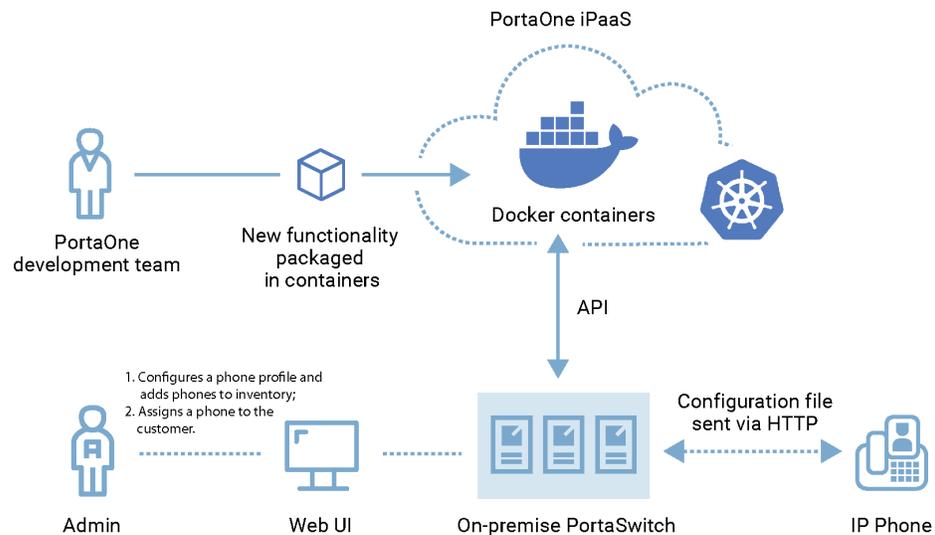
```
% if ($ENV{HTTP_HOST} =~ m/supercall.com/) {  
  
% } else { # our default domain  
  
% }
```

15. Operations and integrations

PortaOne iPaaS

PortaOne iPaaS (Integration Platform as a Service) is a provisioning solution that enables you to deploy new functionalities (e.g. provision new CPE profiles) through containers that run in the cloud. This eliminates the need to update all of PortaSwitch® to the newest release. Thus, you reduce the time to market for new services.

A container is a portable software unit where all the components of new functionality (code, runtime, system tools, system libraries and settings) are stored as one package. Containers are deployed in the cloud-based PortaOne iPaaS independently from the PortaSwitch® installation. As a result they are easily modified and extended.



To start using new features via the PortaOne iPaaS, subscribe to it. To find out about our subscription terms, contact our [Sales department](#).

Benefits:

Service provisioning via PortaOne iPaaS:

- Significantly saves your time and resources since you do not need to update your installation to receive a new functionality.
- Reduces time to market for new functionalities, thus makes you more competitive.
- Improves the customer experience. You can provide your customers with the newest devices and functionalities in the short term.

CPE provisioning via PortaOne iPaaS

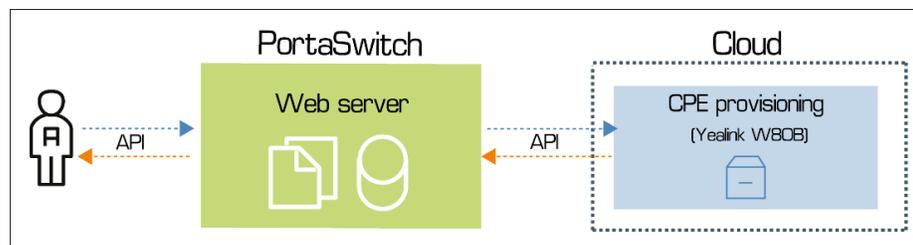
The first device that's available for provisioning via PortaOne iPaaS is Yealink W80B – an IP multi-cell base station. To read about other IP phones provisioning, refer to the **Customer Premises Equipment (CPE) provisioning and inventory** chapter.

This is how CPE provisioning via PortaOne iPaaS works:

Let's say the service provider, Easy Call, wants to provision Yealink W80B to their customers. Since this device model is available via PortaOne iPaaS, Easy Call subscribes to the CPE provisioning service and enables it on their installation.

The rest of the configuration process for creating a new CPE profile and managing the IP phone inventory in PortaBilling® remains the same as before. Refer to the **Auto-provisioning of IP Phones** handbook for the detailed configuration.

When an administrator creates a CPE profile for Yealink W80B, PortaBilling® queries the CPE provisioning container via the API to generate a CPE profile. The CPE provisioning container returns the generated CPE profile and it is stored in the PortaBilling® database along with the internal CPE profiles.



Then the administrator creates the Yealink W80B CPE device and assigns it to an account. PortaBilling® generates a configuration file with the CPE device parameters (username, password, device's MAC address, etc.) and sends it to the CPE provisioning container. This configuration file is used in the CPE provisioning container to adjust Yealink W80B's internal configuration.

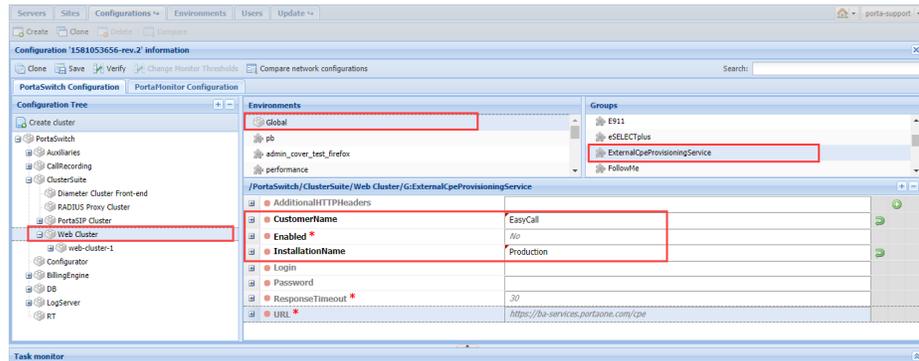
Now Yealink W80B is ready to use by end users.

Configuration:

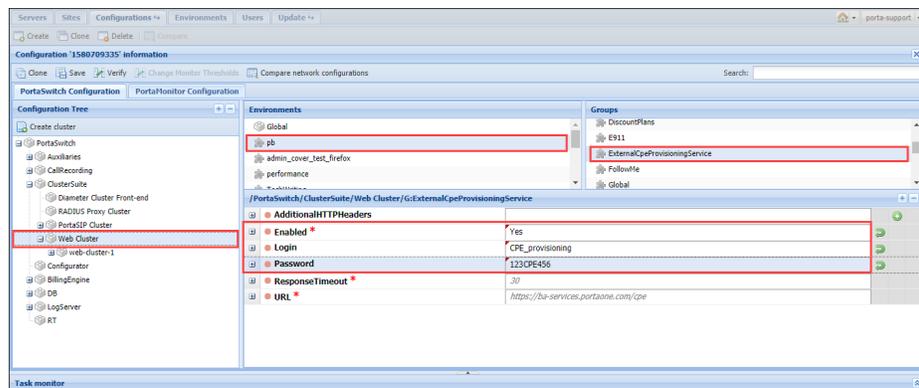
Once you subscribe to PortaOne iPaaS, PortaOne support provides you with credentials to enable CPE provisioning on the Configuration server.

- Go to the **Web_Cluster.Global.ExternalCpeProvisioningService** group

and specify the **CustomerName** and **InstallationName** options as defined in RT.



- Go to a specific environment and set the **ExternalCpeProvisioningService.Enabled** option to **Yes**.
- Specify the credentials provided by our support team under the **Login** and **Password** options.



Integration with the Microsoft Active Directory

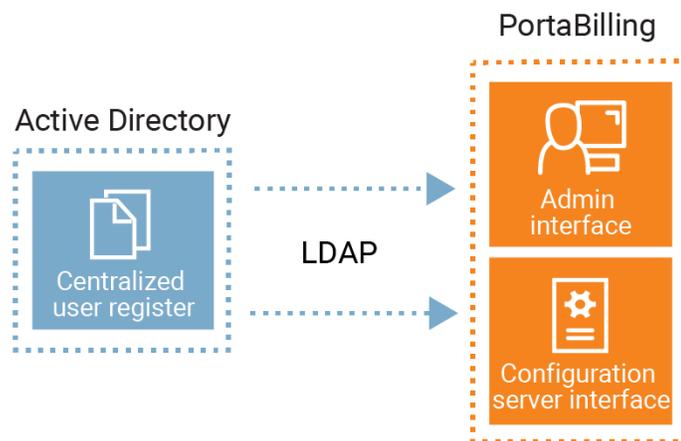
The Microsoft Active Directory (AD) is a directory service that enables service providers to manage users within their entire infrastructure from a single location. An AD administrator can create/block/delete user records, manage their details, and provide access to all corporate resources (for example, Microsoft SharePoint, Microsoft Teams, or Odoo CRM) via the AD. Thus, an AD administrator can quickly enable access for new employees and instantly block access to all corporate resources when somebody leaves the company. This saves significant time on user management and it improves infrastructure security.

An AD administrator can create user records in AD to manage access to:

- PortaBilling® web interface
- The Configuration server web interface

Therefore, users can access the PortaBilling® and Configuration server's web interface in addition to other resources using a single set of credentials. This simplifies the login process and reduces the risk of losing or forgetting multiple passwords.

To communicate with the AD, PortaBilling® supports the Lightweight Directory Access Protocol (LDAP). The (LDAP) user password is stored in the AD database. Thus, the user is authenticated in PortaBilling® through the AD. The LDAP is also used for data synchronization (user name, login, password, email, country, role, expiration date) between the AD and PortaBilling®.



This is how it works:

1. The AD administrator creates a user record in the AD that is then automatically created in PortaBilling®.
2. The user logs in to the PortaBilling® web interface using their LDAP credentials provided by an AD administrator.
3. PortaBilling® sends an authentication request to the AD via the LDAP.
4. The AD verifies the user login and the LDAP password.
5. Access to the PortaBilling® web interface is provided according to the user access role.
6. When an AD administrator changes a user's access role or user details, the data is synchronized with PortaBilling® via the LDAP.
7. When the AD administrator deletes the PortaBilling® user record from the AD, it is automatically deleted from PortaBilling®, too.

Let's say a service provider hires a new employee, David. David is an accountant and needs read-only access to PortaBilling®. The AD administrator opens the AD and creates a new user record with the Helpdesk access role. The AD sends the user data via the LDAP and creates a new user record in PortaBilling®. Now David can log in to PortaBilling® using his LDAP password. In a few months, when David leaves the company, the AD administrator blocks David's user record in the AD. The AD then synchronizes with PortaBilling® and blocks the user record there. Thus, David's access to PortaBilling® is blocked.

Specifics:

- To log in to PortaBilling® or the Configuration server web interface, the AD administrator must create a user record in the AD first.
- The AD stores the user data (e.g., login, password) and authenticates the user's access to PortaBilling® via the LDAP.
- Only an AD administrator can create/delete/block the PortaBilling® user record and manage the mandatory user fields:
 - Name / Surname
 - Email
 - Login
 - Password
 - Role
 - Country
 - Company
 - Expiration date
- User details specified in the AD are read-only in PortaBilling®.
- By default, data in PortaBilling® synchronizes every 10 minutes with AD. You can change synchronization frequency on the Configuration server.
- Access to the Configuration server via SSH is disabled.

Integration with the Microsoft Active Directory allows service providers:

- to enable users to access their corporate resources with a single set of credentials, and
- to manage PortaBilling® users within a centralized user register. This reduces significant time on user management for company administrators.

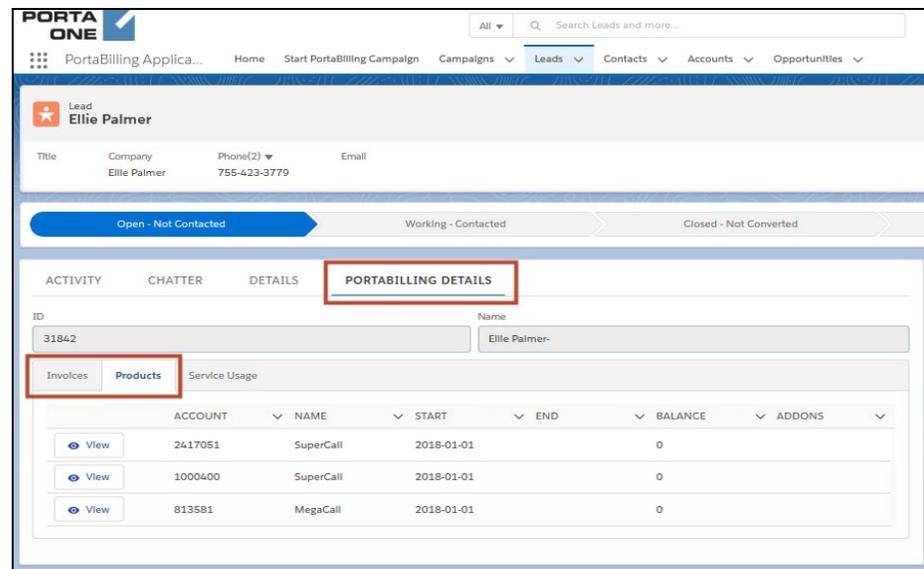
Contact the PortaOne support team to configure PortaBilling® to interact with the Microsoft Active Directory.

Integration with Salesforce CRM

PortaBilling® is integrated with the Salesforce customer relationship management (CRM) system which is designed for the automation of business processes. This integration enables your sales team to:

- perform marketing campaigns for PortaBilling® customers via CRM; and
- use all the CRM features that are built into the system to: contact and follow up with customers, review history, analyze campaigns, etc.

To create a marketing campaign, your sales manager needs a list of customers that he can contact (e.g. via call or SMS). He filters the existing PortaBilling® customers into different groups by using different criteria – country, product, etc. Selected customers become campaign members. When your sales manager opens a customer profile in Salesforce, the **PortaBilling details** tab appears and there he can see the customer’s payment history and products.



For your sales team’s convenience, they can do the following, via the CRM, without logging into PortaBilling®:

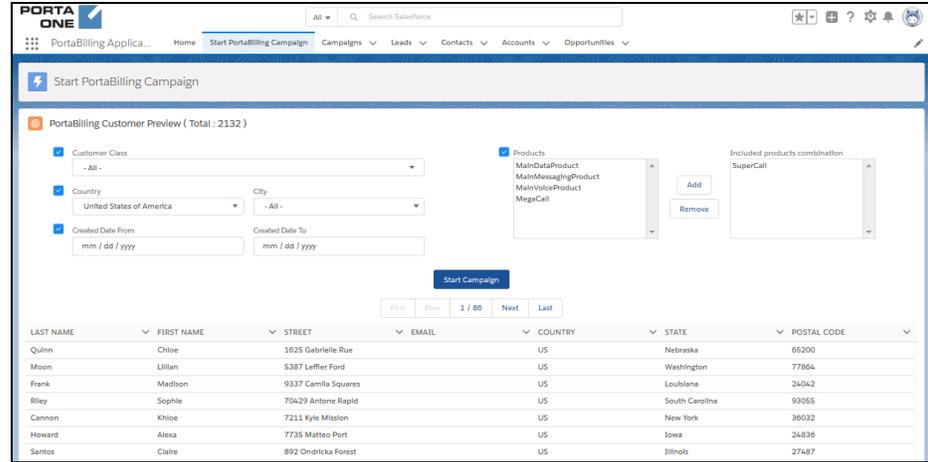
- access relevant information about customers (e.g. their products, invoices and balance), and
- manage customer products and / or add-on products.

Consider the following example:

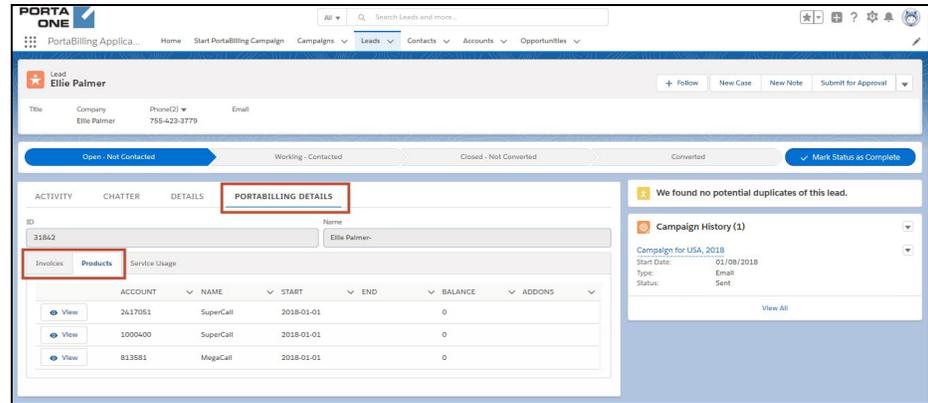
A sales manager, John Doe, creates a marketing campaign in which he wants to sell international calls to Mexico. This is already configured in

the MexicoGo add-on product. So to make this happen, he filters his customers for those who:

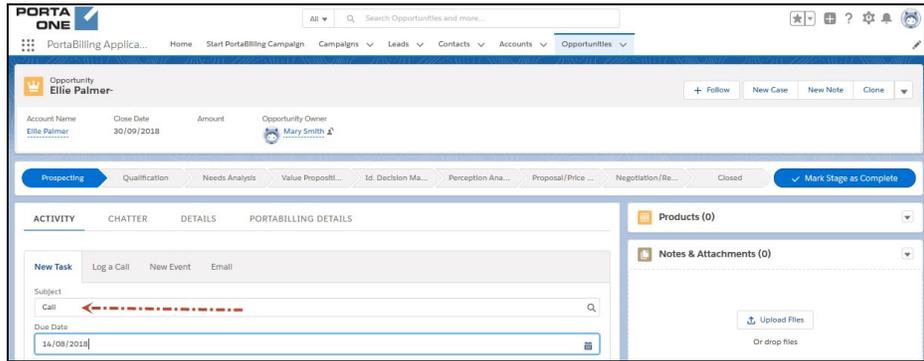
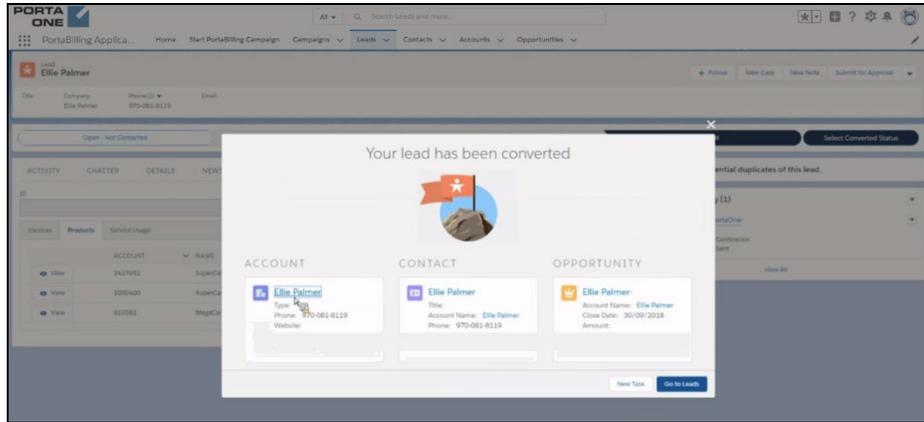
- live in the USA;
- use the SuperCall main product.



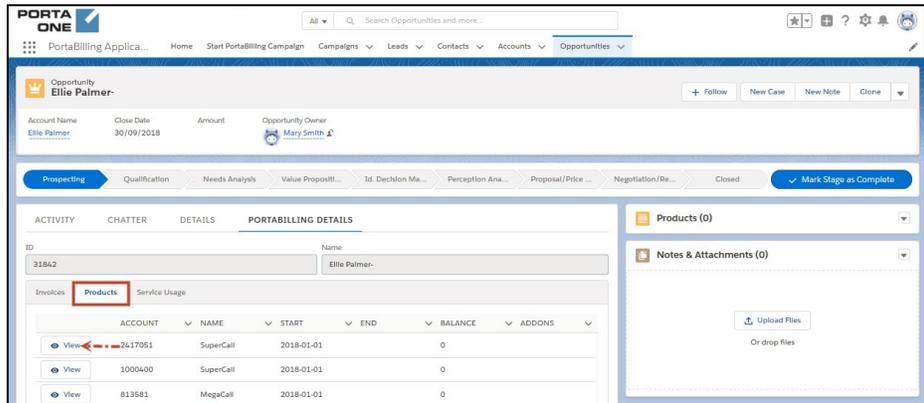
Once the list of customers is ready, John opens the first customer profile, that of Ellie Palmer. Aside from the standard dashboards (e.g. to log calls, create tasks, etc.), there is a **PortaBilling details** tab that gives complete information regarding Ellie’s invoices and products.

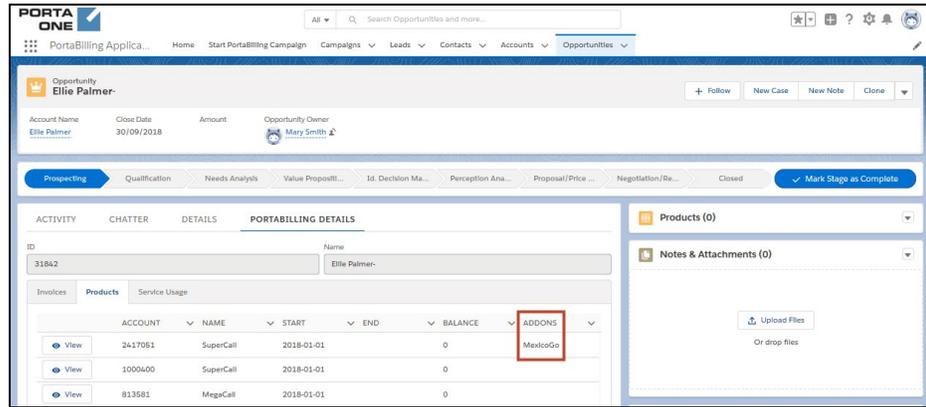
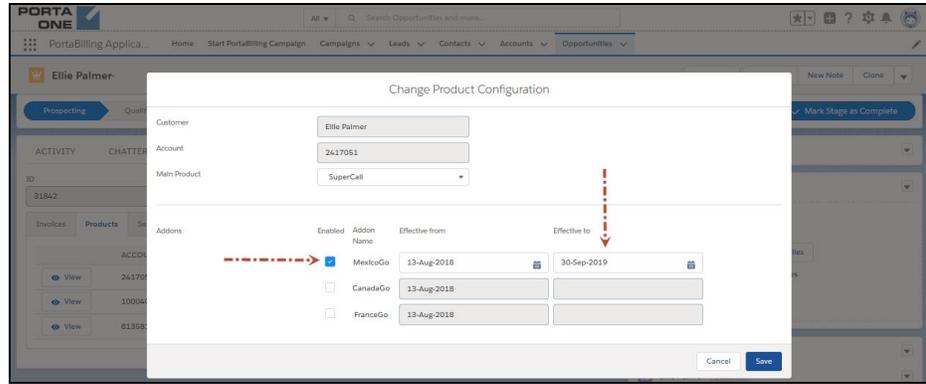


John calls her and offers additional services – i.e. 100 minutes of free calls to Mexico (the MexicoGo add-on product). Ellie Palmer is interested and asks that she be called back the following day to discuss the offer in detail. John then converts Ellie Palmer from the Lead category to the Opportunity, Account and Contact categories (the standard Salesforce sales process) and creates a task to remind him to call her the following day.



The following day, after hearing about the offer in more detail, Ellie Palmer agrees to purchase the new service. To update her service package, John clicks the **PortaBilling details** tab. On the **Products** section he clicks **View** next to the mandatory Account ID. Then he selects the checkbox in front of the **MexicoGo** add-on product and adds the required info into the **Effective to** field. Once John saves his changes the CRM relays this updated data to PortaBilling®. John then closes the deal for Ellie Palmer.





You can find more information about the Salesforce CRM on the [website](#). Please contact the PortaOne® support team for assistance with how to configure the Salesforce CRM.

This integration helps boost your sales by organizing your marketing campaigns. In addition, it allows your sales team to manage customer products via the CRM web interface, to ensure that they have up-to-date information about your customers’ service usage.

Invoice data export to SAP Business One ERP

PortaBilling® is integrated with SAP Business One ERP (Enterprise Resource Planning). PortaBilling® operates as a product inventory and the OCS (Online Charging System). SAP Business One is used to manage accounting processes (e.g. track accounts payable and account receivable, produce financial reports, etc.). So now you can use both systems in sync to streamline your business processes.

You create customers via a SAP Business One web interface. Their data is then automatically replicated to PortaBilling®.

This integration is to meet the business demands of a Communication Service Provider (CSP) from South Africa. There they use ledger codes to categorize revenue data by services. For example, a unique ledger code is allocated to a particular call direction:

- 400000001 – for local calls.
- 400000002 – for international calls.
- 400000003 – for wholesale calls.

These codes are specified in PortaBilling® for destination groups, subscriptions and manual transactions. First PortaBilling® calculates the charges and groups the data according to the ledger codes. Then PortaBilling® creates and sends the invoice report to SAP Business One via the SOAP API. The resulting report contains the accumulated costs for specific services and customers all grouped by ledger codes.

An accounting manager can now produce general ledger reports via SAP Business One for a specific period, e.g. an expense report for a particular quarter.

PortaBilling® sends the invoice report to SAP Business One when the invoice is produced upon:

- the customers' billing period closure,
- as an out-of-turn invoice generation, and /or for
- an xDR re-rating and an invoice regeneration.

To send data from PortaBilling® to SAP Business One, you must use middleware. PortaBilling® must communicate with that middleware via the SOAP API.



Since your business requirements most likely differ from South Africa CSPs, you need to develop your own middleware.

This integration automates the data transfer for invoicing and reporting to SAP Business One.

Troubleshooting Tools

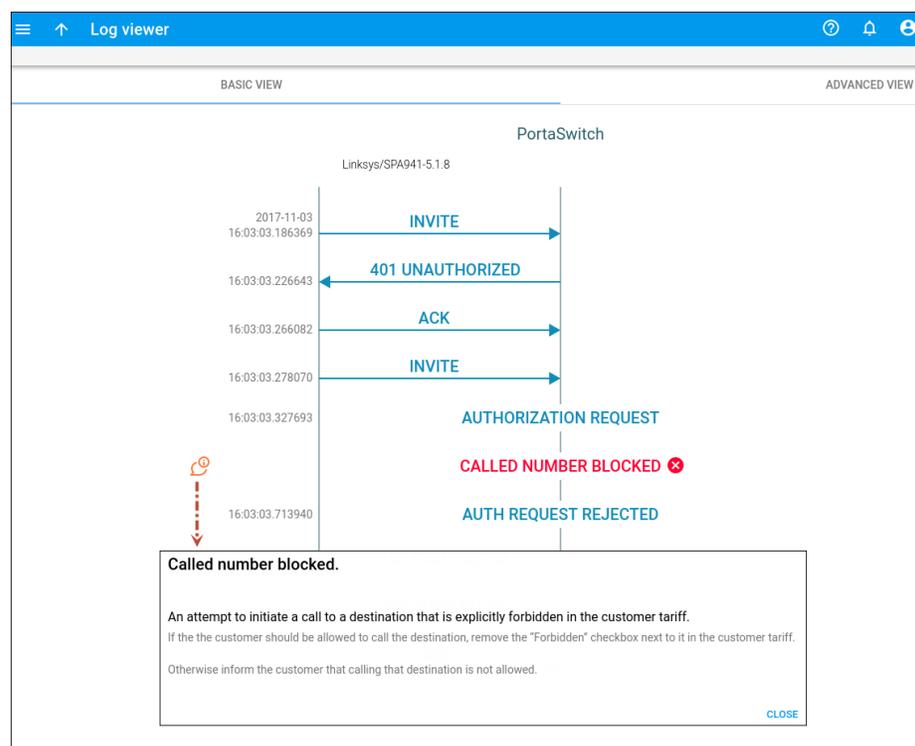
Trace session functionality is an efficient tool that allows user sessions to be traced and simplifies troubleshooting for administrators.

The key features of trace session functionality are:

- A single page for session search and display;

- Merged SIP and billing logs to provide an entire overview of the session flow;
- General information about a session in the form of a diagram between the main participants (e.g. SIP phones and PortaSwitch®) with the issue description and troubleshooting instructions for the most common issues; and
- The ability to switch to the advanced log view if deeper investigation is required (e.g. to review the call flow among nodes within the PortaSwitch® cluster).

For instance, on a call with an administrator, John Doe (12103355787) wants to find out why he cannot call 18660333777. The administrator opens the trace session page and searches the session using the account number and time interval. On the log viewer page, he sees that John Doe is not allowed to call the destination because it is forbidden.



For backward compatibility, an administrator can access billing logs on the old web interface by clicking the **Old UI** icon.

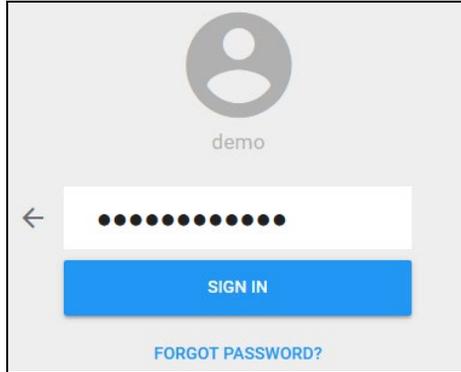
This feature drastically simplifies troubleshooting and reduces the load on your administrative staff.

Please refer to the Troubleshooting section of the [Unified PortaSwitch Handbook Collection](#) for more troubleshooting tips.

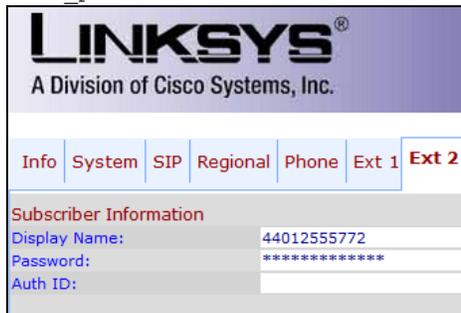
Password protection

There are two types of passwords in PortaBilling®:

- **Web interface password.** This password is used in combination with a login to authenticate a user (e.g. admin, end user) who gains access to their web self-care interface.



- **Service password.** This password is used to authenticate all calls made using the account. It is also used to register a SIP account on a SIP phone / softphone. Only accounts have this password. In the PortaBilling® API, a service password corresponds to a `h323_password` attribute.



To secure password storage, hashing and encrypting mechanisms are used. These measures protect users' passwords in case an unauthorized person obtains access to the database.

Hashing passwords

Hashing performs a one-way password transformation. A password is turned into a hash value (data chunk of a fixed size). This hash value is unique and stored in a database. When a user attempts to log in, the system takes the password entered and performs a similar one-way hash. Then the system compares the hash value that was provided with the database hash value. If they match, the user logs in successfully.

It is impossible to transform a hash value into initial symbols. Hashing is not reversible.

Highlights:

- Passwords are hashed for:
 - All users that are created via the PortaBilling® web-interface (My company > Access control > Users), e.g. users with root permissions, administrators, operators, etc.
 - CC staff created by resellers. For example, customer care support employees.
- Web interface passwords are hashed only.
- Passwords are hashed by default. Hashing cannot be turned off.
- Hashed passwords cannot be converted to plain text.

Encrypting passwords

Encryption is a two-way password transformation. It means that a password can be decrypted to initial symbols. Users with ACL / a role that permits them to obtain passwords, can obtain them in plain text. For example, an administrator can put an account password into plain text by using the `get_account_info` API method.

Encryption uses a passphrase along with salt. The passphrase is a key used for encrypting or decrypting passwords. Salt is random data used to randomize output after encryption. Salt complicates the process of deciphering an encrypted password since each salt value is unique, even for two identical passwords.

Passwords encrypted with salt are stored in a database while a passphrase is stored on several servers. Encrypted passwords and a passphrase are stored in different places which makes them secure. For this reason, it is difficult to crack these passwords, even if an attacker obtains access to the database.

An administrator can change a passphrase with a CLI utility. For more information, contact PortaOne support. During this procedure, an old encryption passphrase is replaced by a new one. This requires decryption-encryption and an update for all passwords in the database. During this procedure, account owners can continue using their services as usual; i.e. there is no downtime.

Highlights:

- Passwords are encrypted for the following entities:
 - Accounts
 - Customers
 - Resellers
 - Distributors
 - Representatives

- Both web interface passwords and service passwords are encrypted.
- The passwords of vendors are always stored in plain text.
- Encryption is enabled by default. It can be turned off on the Configuration server (**Security.EncryptPasswords** option).
- A passphrase can be changed via a console utility.
- Users with the right ACL / role permissions can obtain encrypted passwords in plain text.

Password expiration

By default, all web interface passwords that are used to log in expire in 30 days. To set another timeframe, go to the Configuration server (**Web.Password_expire** option).

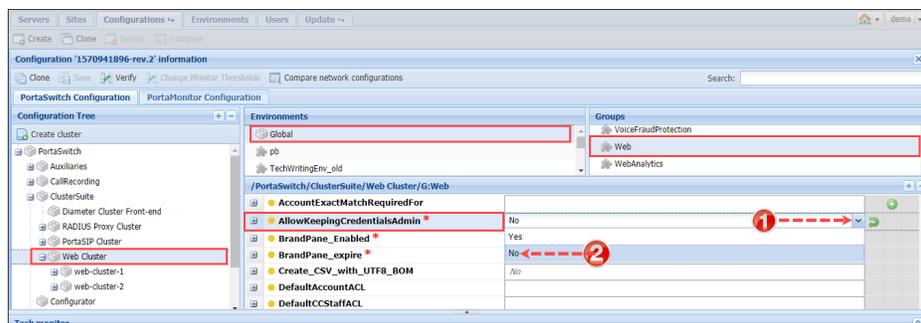
These protection measures increase the security level of users' passwords stored in a database.

Storing admin web password in browsers

Administrators often store credentials to PortaBilling® web GUI in password managers. The best security practice is to store the passwords encrypted. However, a browser's password manager does not encrypt passwords by default and those passwords might be stolen (e.g. by means of malicious software, when the device is lost, etc.).

To ensure the optimum security in PortaBilling®, you can disable storing credentials in browsers' password managers. Thus, every administrative user must specify their login and password every time they log in to PortaBilling®.

To disable storing credentials, set **AllowKeepingCredentialsAdmin** option to **No** on the Configuration server.



16. Statistics and reporting

Date and time information

Time zones

Due to the global nature of the VoIP business, it is very important to identify times with reference to time zones. For example, there is an 18-hour difference between 8 p.m. Melbourne, Australia, and 8 p.m. Seattle, USA – which could mean many thousands of calls and thousands of dollars incorrectly counted if the time zone is not taken into consideration.

This is why every object in PortaBilling® (e.g. user, customer, vendor or account) has a time zone associated with it. All dates are stored in the database in a universal, portable representation, and are converted into the required time zone when necessary. For example, say a call is made at 14:03 New York time (your local time zone). The call is made using a calling card running on the America/Chicago time zone, and the account belongs to customer Easy-Cards-Reseller, whose time zone is America/Los_Angeles. The call goes to vendor Deutsche Telekom, which has Europe/Berlin as its time zone.

Thus, if each of these entities views information about the call, they will see it in their respective time zone, as follows:

- Account: 13:03
- Customer: 11:03
- Vendor: 20:03

This helps to present time information in a format suitable for the user. Every customer also has a billing time zone that affects the billing period. Thus if a customer has billing time zone Europe/Berlin and a monthly billing period, this means that his billing period covers one whole month, from 00:00 on the first day of the month until 00:00 on the first day of the next month **in his billing time zone** – Western European Time, in this example.

Also note that all time zones are expressed in relation to a geographical location in PortaBilling®, for example, Europe/Prague. This helps to avoid ambiguous abbreviations such as EST, which could mean both New York time and Melbourne time. Moreover, by using such a notation daylight savings time automatically comes into effect when applicable, so you do not have to think about whether it is EST (Eastern Standard Time) or EDT (Eastern Daylight Savings Time).

Date and time formats

Times and dates have different formats in different places in the world. For example, the same date will be represented as 20.05.2003 in Germany and 05/20/2003 in the US. It is important that the user be able to work with dates in his native format in order to facilitate work and avoid data interpretation problems.

Each PortaBilling® user (admin, account, customer care) has his own date and time formats. When logging in to the web interface, all date and time information will be shown in the user's specified format. Additionally, all data input fields will also assume the date and time formats specified for the given user.

xDR summary

An administrator can browse summary information for all xDRs for a selected date and time range. You can browse the xDR summary for customers, accounts (credit and debit), vendors, resellers and distributors.

The xDR summary includes the following information:

- Total number of xDRs filtered;
- Totals by services;
- Total sum of payments and credits;
- Total sum of taxes.

xDR summary		
for the period 2019-05-11 05:35 - 2019-06-11 05:37		
Total xDRs number: 14		
Voice Calls, min	9.76667	
Charges: 0.10167 USD		
Service usage, USD	Recurring charges, USD	Other charges, USD
0.10167	0.00000	0.00000
Payments and credits: 0.00000 USD		
Payments, USD	Refunds, USD	Credits & promotions, USD
0.00000	0.00000	0.00000
Taxes: 0.00000 USD		

An administrator can download the xDRs for a specific period for further analysis, statistics purposes, etc. The output file is in the .csv format.

The data in the xDR file differs depending on the type of entity (e.g. customer, vendor) and the selected service. An administrator can select a

specific service (e.g. voice call) or download xDRs for all available services including subscriptions, payments, taxes, DID usage, etc.

Customer xDRs

This is an example of the output file with the customer xDRs for voice calls service:

Account	From	To	Country	Description	Last hunting group	Last extension	Connect time	Disconnect time	Charged time, hour:min:sec	Amount, USD	Status
1206550017	1206550021	1204257833	United States of America				5/12/2019 14:14	5/12/2019 14:14	0:00:06	0.01	
12024257833	12024257833	1866747647	United States of America	North America Toll-Free			5/14/2019 9:50	5/14/2019 9:52	0:02:42	0.27	reverted
1206550011	1206550011	1206550023		Calls between IP phones			6/11/2019 14:59	6/11/2019 14:59	0:00:06	0	
1206550011	1206550011	1206550023		Calls between IP phones			6/11/2019 14:59	6/11/2019 14:59	0:00:26	0	

The following data is available:

- Account ID
- From (CLI or ANI)
- To (CLD or DNIS)
- Country
- Description
- Last hunting group
- Last extension
- Connect time
- Disconnect time
- Charged time (in hours:minutes:seconds)
- Amount (USD)
- Status

NOTE: The call duration shown in the file is based on the charged duration of the call, not the actual call duration. So if, for instance, you bill by 6-second intervals, the total call duration will be higher than the actual duration of the calls.

Vendor xDRs

It is very useful to have a detailed list of xDRs for a specific vendor in case of disputes over the amount of terminated traffic.

This is an example of the output file with vendor xDRs for voice calls service:

Account	From	To	Country	Description	Connect time	Disconnect time	Disconnect Reason	Charged time, hour:min:sec	Amount, USD	Status
12045854745	12123344567	12045854745		Default	4/2/2019 8:48	4/2/2019 8:48	Normal call clearing	0:00:24	0.4	
12045854745	12123344567	12045854745		Default	4/2/2019 8:49	4/2/2019 8:49	Normal call clearing	0:00:20	0.33	
12123344567	12123344567	12045854745		Default	4/2/2019 8:50	4/2/2019 8:50	Normal call clearing	0:00:07	0.11	
12123344567	12123344567	12045854745		Default	4/2/2019 8:50	4/2/2019 8:50	Normal call clearing	0:15:08	15.13	
12123344567	12123344567	12045854745		Default	4/2/2019 8:51	4/2/2019 8:51	Normal call clearing	0:00:13	0.22	
12123344567	12123344567	12045854745		Default	4/2/2019 8:53	4/2/2019 8:53	Normal call clearing	0:27:04	27.06	
12045854745	12123344567	12045854745		Default	4/2/2019 8:53	4/2/2019 8:53	Normal call clearing	0:00:32	0.53	
12045854745	12123344567	1206550025		Default	4/2/2019 8:53	4/2/2019 8:53	Normal call clearing	0:00:26	0.43	
12045854745	12123344567	12045854745		Default	4/2/2019 8:53	4/2/2019 8:53	Normal call clearing	0:01:03	1.05	
12045854745	12123344567	12045854745		Default	4/2/2019 8:54	4/2/2019 8:55	Normal call clearing	0:01:12	1.2	

The following data is available:

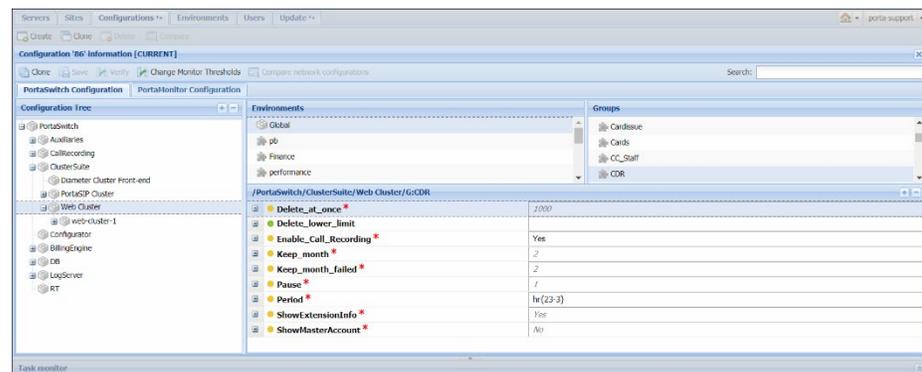
- Account ID

- From (CLI or ANI)
- To (CLD or DNIS)
- Country
- Description
- Connect time
- Disconnect time
- Disconnect reason
- Charged time (in hours:minutes:seconds)
- Amount (USD)
- Status

xDR cleanup procedure

xDRs are stored in the Main database and replicated to the Replica database. In order to prevent system overgrowth, xDRs are periodically cleaned up. This means that a special “Cleanup” xDR exists which accumulates charged amounts and times from all xDRs that are cleaned up for each account / customer / vendor. This xDR shows up as a “Cleaned up transaction” on the xDR history page of the web interface.

By default, the Cleanup procedure occurs during each off-peak period (every night) for all xDRs that are more than 2 months old, though these options are configurable to meet special requirements via the Configurator.



Parameter	Value
Delete_at_once	2000
Delete_lower_limit	
Enable_Call_Recording	Yes
Keep_month	2
Keep_month_failed	2
Pause	1
Period	hr(23-3)
ShowExtensionInfo	Yes
ShowMasterAccount	No



Although all the statistics / invoices are calculated up to the moment of the specified xDRs’ Cleanup, you can increase the period for keeping old xDRs according to your needs.

Custom reports

If your business process requires that specific information from PortaBilling® be presented in a certain format (e.g. your marketing department needs to know the top 10 countries based solely on calls longer than 15 minutes, including the ratio between these calls and all

other calls for each country), and this information is not available among PortaBilling® default reports, then there are two options available:

- Create your own custom report using tools such as Crystal Reports (for more details, see the **Custom Reports** section). This assumes that you have sufficient experience with relational databases and reporting tools.
- Contact the PortaOne team, so that a custom report can be developed for you and uploaded to your system. In this case, no technical knowledge on your side is required.

PortaBilling® custom report module allows the creation of “report” objects, which can later be uploaded to your system and then executed by your administrators. In this way, you can maintain an unlimited number of different reports specific to your business.

Report type

Report type defines the purpose of a report: does it show the most popular destinations, or the customers who make the most calls? Several types of reports are already supplied with PortaBilling®, including:

- Most popular destinations of a specific customer.
- Total amount of minutes and charges for a specific customer.

Other types of reports can be developed on request. When a new report object is uploaded to your system, a new type of report will appear.

Report type defines:

- What types of input parameters are required by the report (e.g. customer name, time interval, etc.).
- What the output values of this report are. Since the report results are displayed as a table, this defines what the columns represent in this table, while every row in the table represents another line of report data.

Report query

Whereas report type defines a report’s general characteristics (e.g. “Most popular destinations”), by executing a report query you may obtain a specific result (e.g. “What were the most popular destinations for customer John in the last two weeks?”).

Sometimes you will simply execute a query once (for instance, you would like to quickly verify that a particular customer did indeed make more than 1,000 calls yesterday), while other types of queries will be made routinely (e.g. every morning your sales manager checks the previous day’s statistics for his VIP customers). To accommodate both types of needs, PortaBilling® allows you to create a new query and then:

- Run it once in your web browser to see the results immediately.

- Save the query (including input parameters) for later use. When you need to re-run the query sometime in the future, you simply open and execute it, i.e. you do not have to fill in any parameters. This significantly reduces the time required to prepare complicated reports, and allows you to minimize the chance of error. The administrator creates the query just once, and then inexperienced users need only execute it.

Also, to save you time when preparing multiple similar reports, PortaBilling® allows you to create a query as a copy of another existing query. After you have defined a complex report with many input parameters for customer A, simply create a copy of it and change the customer name to “B”. Now you have reports for both customers, either of which can be executed with just one mouse click.

Query execution

There are several methods for executing a query and obtaining report results:

- You may request query execution from the PortaBilling® admin interface, and see the result in your web browser.
- You may request that a query be executed in the background, with the results delivered to you via email. This is the preferred method for complicated reports involving large arrays of data, since in this case PortaBilling® can schedule query execution in such a way that it does not interfere with other tasks. This allows you to avoid performance degradation when several users are trying to execute reports concurrently; PortaBilling® will execute all these reports sequentially.



This is the recommended method for executing reports.

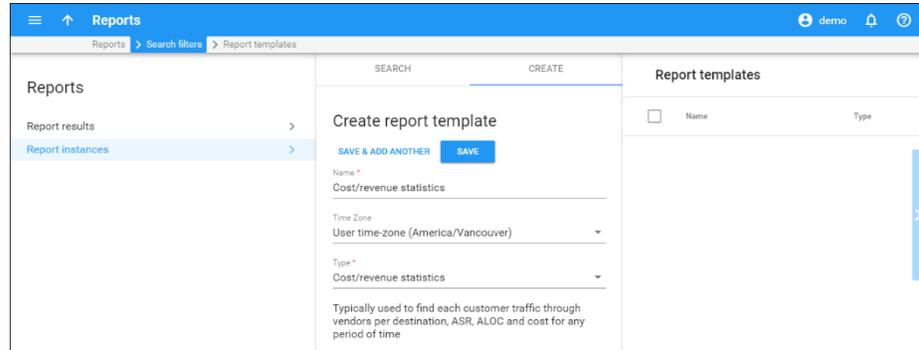
- You may schedule a query to be executed at a certain moment in time or periodically, with the report results delivered to you via email. This allows you to automate reports that you use on a continuous basis. For instance, you may schedule a “Total minutes per customer” report for 3 a.m. every day; then every time you open your mailbox in the morning, you will find all the required information. You can view a list of scheduled query executions and their status on the PortaBilling® admin interface.

Query results

When you execute a report query from the web interface, its result is stored in PortaBilling®. Thus if you did not save the results initially, you need not run the report again, as the data for the last ten report executions is always available.

Custom reports supplied by default

PortaBilling® includes a comprehensive set of reports, which are already included in the system upon initial installation:



Report	Purpose
Accounts receivable with aging	Shows the sum of unpaid customer invoices (accounts receivable) broken down by date ranges of 0-30, 31-60, 61-90 and >90 days. Administrators can obtain accounts receivable figures as of the moment of report generation or as of any date in the past.
Accumulated costs by destinations	Shows the accumulated costs of a service for a specific vendor, grouped by destination.
CDRs for all sub-customers of a reseller	Provides information about CDRs produced for all subcustomers of a reseller within a certain period. With this report, administrators have comprehensive information about charges applied to all reseller's subcustomers in one place. Resellers, in turn, may use this report for billing purposes. For example, they can process the output file using a third-party accounting system for generating invoices or various statistical reports.
Charges summary	Allows you to obtain a monthly report about a specific reseller's customers and their charged amounts in the form of a single list (.csv file format). This is convenient for subsequent computer processing (typically by a Reseller). Its main purpose is to free the Reseller from having to conduct endless customer searches.
Cost / Revenue Statistics	Provides a detailed information about the profitability of specific customer and vendors. Also shows ASR, ALOC, cost and other parameters for any given period of time.

Cost / revenue statistics advanced	Similar to the above, but this allows you to group data using the customer name and account billing mode or show details for a specific customer site. It also shows additional customer information (e.g. tax codes, contract numbers, etc.) and allows that to be used for grouping the data in the report.
Customer completed summary	<p>Provides information about the number of successful and failed calls made over any specified period of time, their total duration and the amount charged to customers for them (in the base currency of the virtual billing environment).</p> <p>You can filter the report data by specifying the following parameters: a service, a customer, a customer class, a representative and a period of time. Please note that you won't be able to choose a customer if you already specified a customer class or representative. However, in this case, the report will display separate rows for every customer under the chosen customer class or representative.</p> <p>This report presents information about calls that have the following disconnect causes: 16 ("Normal call clearing"), 31 ("Normal, unspecified") for successful calls; 17 ("User busy"), 18 ("No user responding") and 19 ("No answer from user") for failed calls.</p>
Customer statement	Prints out a list of a customer's invoices / payments, showing the running balance for a given time interval and the totals for all these categories.
Customer statement advanced	Prints out a list of changes of customer's balance for a given time interval.
General cost/revenue statistics	Shows general cost / revenue statistics for a specific period of time and sorts the data depending upon a selected report type (e.g. by customer and destination, subtotal per country).
Generated revenue by destination	Shows generated revenue for specific services and customers, grouped by destination.
Most popular countries	Allows you to monitor which countries particular customers call most often.
Most popular destinations	Allows you to monitor which destinations (individual prefixes) particular customers call most often.

Payments	Provides details about payments recorded in the system during a specific time interval (with the option of filtering payments for particular customers).
Representative commission report	Provides details about the representative's commission (amounts earned by representative) and the commission payments made to the representative.
Revenue	Shows the total revenue (all the charges to customers) during a specific period. Includes the following report types: <ul style="list-style-type: none"> • Revenue by date. • Revenue by charge type. • Revenue by customer class.
Routing info - LCR blending	Analyzes and shows the available routes, calculates the average termination price for each destination, and then uses this as a basis for computing the rate for a given customer. The most important step is to “blend” the prices of different vendors according to a realistic scenario. For more detailed information, see the section <i>Routing Info – LCR Blending</i> in the Unified PortaSwitch Handbook Collection .
Traffic usage trending	Allows you to view selected customers' daily voice minutes per destination group set. This report evaluates historic data to indicate whether a customer's voice traffic has gone up, down or remained constant. When the report is configured and generated you will find the following indicators: <ul style="list-style-type: none"> • Green indicator color demonstrates an increase of Total Minutes over Average Minutes. • Red indicator color indicates a decrease of Total Minutes over Average Minutes. • Blue indicator color indicates that there was no significant change between the Average Minutes and the Total Minutes.
Vendor average success rate	Shows the vendor average success rate for the specific period.

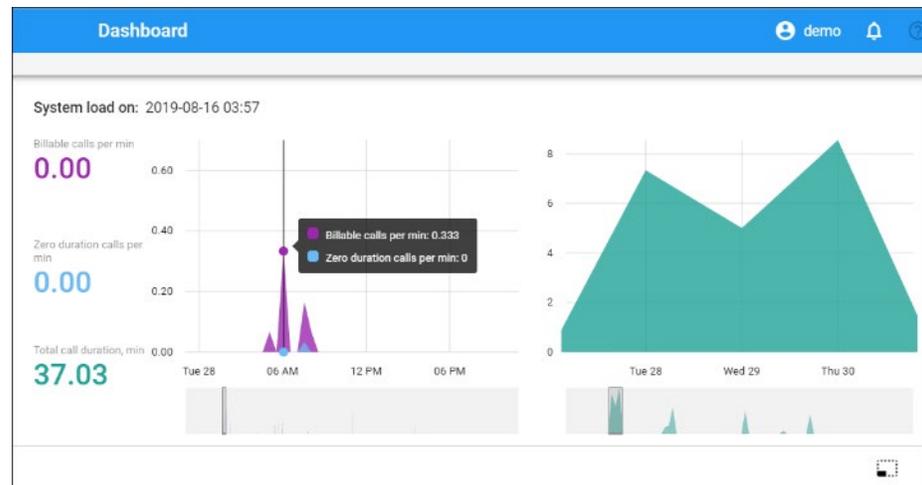
PortaBilling server health monitoring

Entire system load

These statistics are accessible from the **Dashboard** page as a widget.

This widget includes two graphs. The left-hand graph shows the number of calls registered by a production system every 15 minutes: normal calls with duration >0 and calls with zero duration. The right-hand graph shows the total call duration.

The data on both graphs is collected for the last three months. To view the statistics for the particular time interval, select it from the bottom scale.



Number of calls

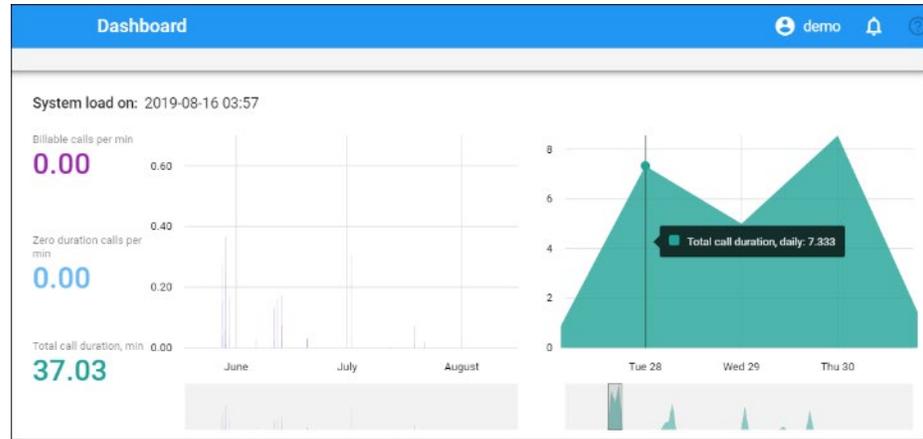
If the disconnect time on the CDR does not fall within the past 15 minutes, the call may not be finished, and hence will not be reflected on the graph.

Zero duration calls

The graph can also show possible problems in the system, such as an unexpectedly high number of failed calls. By default, this graph displays statistics for the last 30 hours.

Total minutes statistics

The graph below gives you the ability to monitor how the call volume in your system changes each day.



VoIP network performance statistics

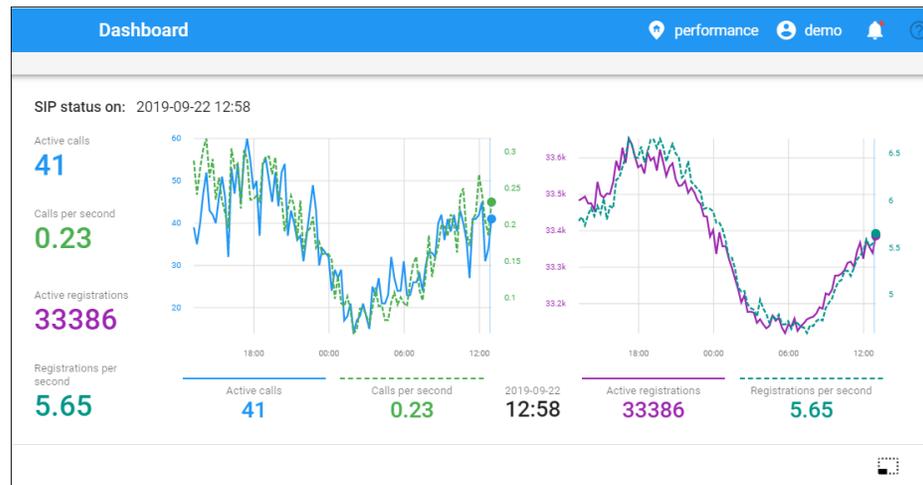
SIP status

You can view statistics about the SIP services you provide for the past 24 hours. This helps you to detect changes in your network (e.g. a burst of registrations) and respond to them accordingly.

The statistics includes information about calls and registrations. There are four metrics:

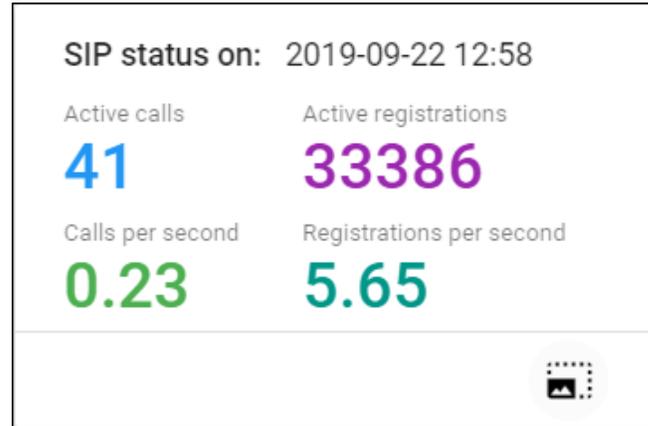
- the number of active calls;
- the number of calls per second;
- the number of registered UAs;
- the number of registrations per second.

They are available as the **SIP status** widget on the **Dashboard** page.



To measure statistics, enable the **Metrics.StorageEnabled** option on the Configuration server. The measurement frequency is 15 minutes. The date and time of the last measurements are shown at the top of the widget.

The SIP status widget has both collapsed and extended views. In the **collapsed** view, only the last measured data is shown.



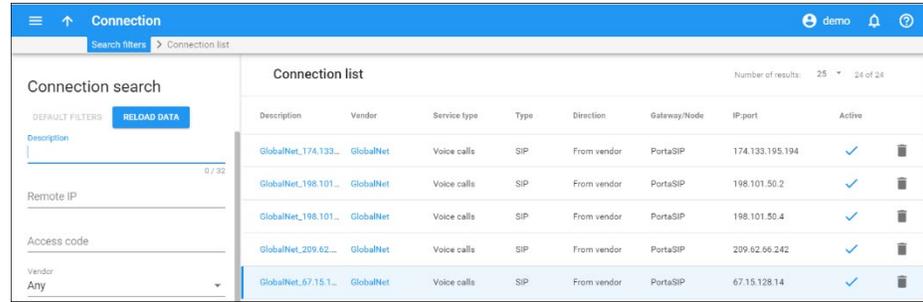
The **extended** view also includes charts where you can see statistics for an exact timestamp during the last 24 hours. There are two charts: the left-hand one shows statistics for calls and the right-hand chart – for registrations.



The figures below the charts show the metrics' values for an exact timestamp. The timestamp is placed between the charts. When you highlight a specific time on one of the charts, the same time is highlighted on the second chart and the values below the charts change.

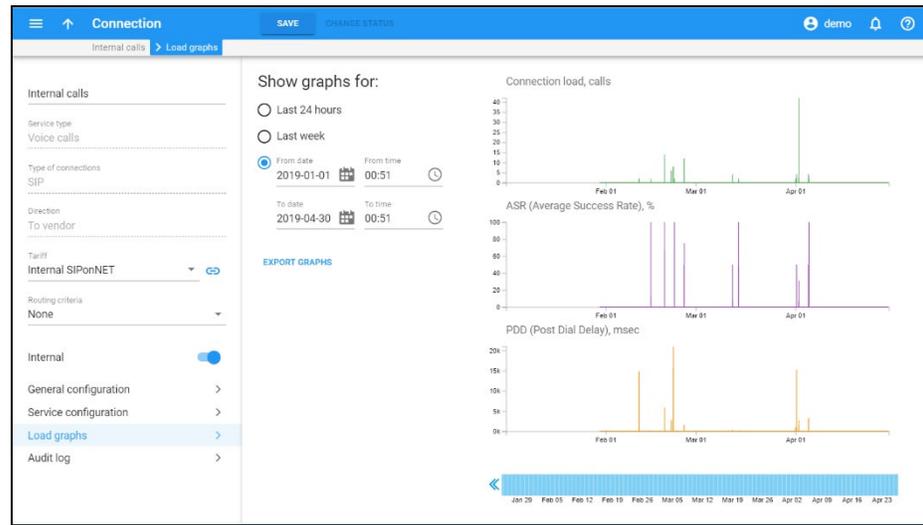
Connection load

You can monitor load and quality parameters for individual connections. To see a list of the available connections, follow the **Connections** link from the **Toolbox** section in the main menu.

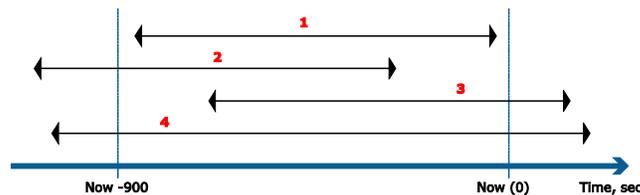


Description	Vendor	Service type	Type	Direction	Gateway/Node	IP:port	Active
GlobalNet_174.133...	GlobalNet	Voice calls	SIP	From vendor	PortaSIP	174.133.195.194	✓
GlobalNet_198.101...	GlobalNet	Voice calls	SIP	From vendor	PortaSIP	198.101.50.2	✓
GlobalNet_198.101...	GlobalNet	Voice calls	SIP	From vendor	PortaSIP	198.101.50.4	✓
GlobalNet_209.62...	GlobalNet	Voice calls	SIP	From vendor	PortaSIP	209.62.66.242	✓
GlobalNet_67.15.1...	GlobalNet	Voice calls	SIP	From vendor	PortaSIP	67.15.128.14	✓

After you select a particular connection, click **Load graphs** to see its load graphs:



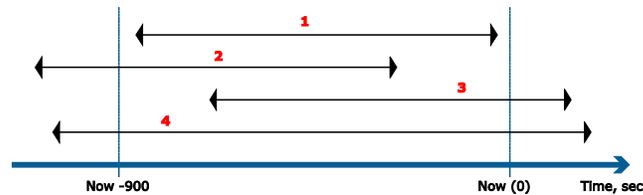
The graphs show the node load, setup time and ASR (Average Success Rate) statistics. Information appears with a one hour delay.



Assumptions: To make load calculations easier, we ignore calls longer than one hour (since these are quite rare, it should not affect the final results much). Also, although a failed call has “zero” duration, it still occupies a port on the gateway for some time (while the gateway is trying to connect the call). Therefore, when calculating, failed calls are regarded as occupying the port for 10 seconds.

Connection load calculations

For a better understanding of the calculations performed, consider the following diagram:



The total duration must be calculated for all calls processed by the system corresponding to the interval between “now - 1 hour 15 minutes”, and “now - 1 hour”.

Four different call types are recognized in the above diagram; an additional call type is “zero duration”:

1. Sum of durations of all the calls.
2. Sum of durations between `t0` and `disconnect_time`.
3. Sum of durations between `connect_time` and `t1`.
4. Calculate sum as number of calls * 900 seconds.
5. Calculate sum as number of calls * 10 sec.

The maximum value in seconds which a connection can hold in the 900 sec. interval is calculated as $\text{capacity} * 900\text{sec}$. The connection load is calculated as the sum of durations for all call types divided by the maximum value in seconds. The result is shown as a percentage value (multiplied by 100).

ASR calculations

Assumption: Connection load less than 5% is not representative for ASR calculating.

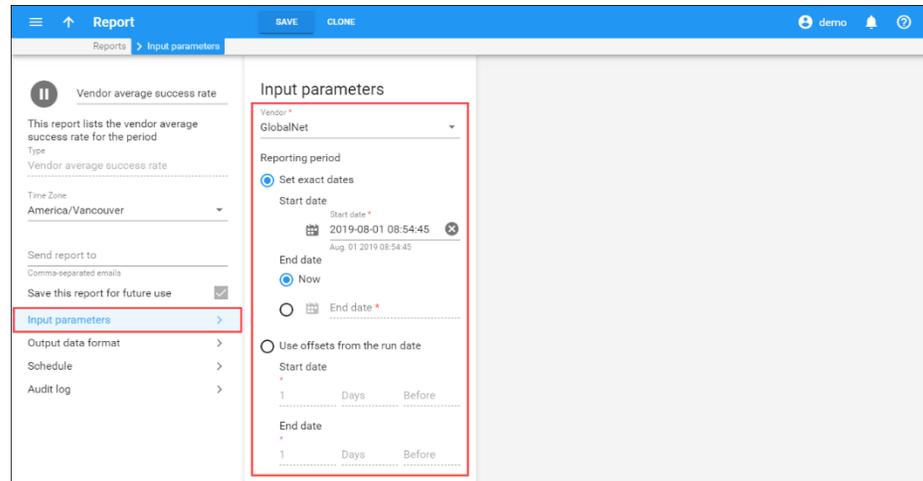
ASR is calculated as the number of connected calls divided by the total number of calls. The result is shown as a percentage value (multiplied by 100).

Relative setup time calculations

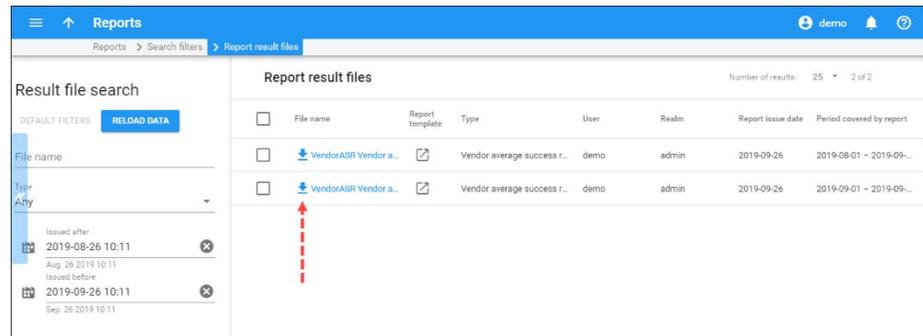
Relative setup time is calculated as the ratio between the total setup time of all calls during the given period and the maximum value in seconds ($\text{capacity} * 900 \text{ sec}$).

Vendor ASR statistics

You can obtain the vendor ASR statistics for a specific time period by executing the **Vendor average success rate** report.



Download the report results in a .csv file.



The ASR statistics look as follows:

TZ	Vendor	From	To						
America/Vancouver	GlobalNet	9/1/2019 8:54	9/25/2019 23:56						
Destination	Destination Description	Country	Service	All Transactions	Successful Transactions	ASR, %	Used Quantity	Unit	ALOC
1800		USA	Voice Calls	40	15	37.5	0:08:11	hour:min:sec	0:00:32
SUBTOTAL				40	15	37.5	0:08:11	hour:min:sec	0:00:32
44		UK	Voice Calls	12	7	58.33	0:01:53	hour:min:sec	0:00:16
SUBTOTAL				12	7	58.33	0:01:53	hour:min:sec	0:00:16
TOTAL				52	22	47.92	0:10:04	hour:min:sec	0:00:24

The data available is:

- Total number of calls.
- Number of calls with a non-zero duration (successful transactions).
- ASR rate.
- Total call duration (used quantity).
- ALOC (Average Length of Call).

The values are aggregated per destination prefix, with a subtotal per country and a total for all calls.

17. ■ Appendices

APPENDIX A. Supported RADIUS RFCs

- RFC 2865 (“Remote Authentication Dial In User Service (RADIUS)”) – supported.
- RFC 2866 (“RADIUS Accounting”) – supported.
- RFC 5090 (“RADIUS Extension for Digest Authentication”) – partially supported.
- RFC 3579 (“RADIUS (Remote Authentication Dial In User Service) Support for Extensible Authentication Protocol (EAP)”) – partially supported.
- RFC 3748 (“Extensible Authentication Protocol (EAP)”) – partially supported.

APPENDIX B. Supported Diameter RFCs

RFC 6733 (“Diameter Base Protocol”) – supported.

RFC 4006 (“Diameter Credit-Control Application”) – supported.

APPENDIX C. Change-of-Authorization (CoA) and Disconnect Message (DM) packets supported by PortaBilling®

The following node types support DM codes (the “40 – Disconnect-Request,” “41 – Disconnect-ACK,” “42 – Disconnect-NAK” specified in RFC3576):

- Cisco
 - NAS
 - ASR
 - PDSN
- MikroTik
- ECI

The following node types support CoA (“43 – CoA-Request,” “44 – CoA-ACK,” “45 – CoA-NAK” specified in RFC3576):

- Huawei
- ECI

To find out whether a particular piece of equipment that you want to connect to PortaBilling® supports DM or CoA or both packets, please refer to the documentation for that equipment.

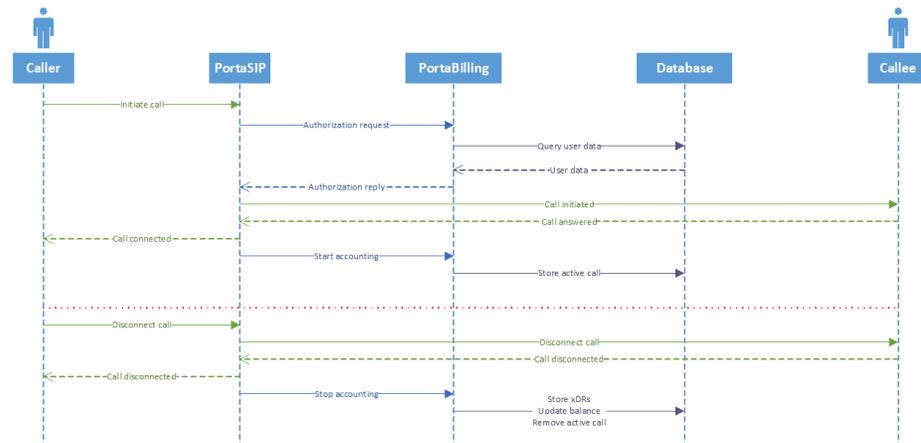
APPENDIX D. The Number_Portability table description

Field	Type	Description
i_number	int	The unique ID of the ported number record in the database
destination	string	The ported number. When a call is being established, PortaBilling® will compare the originally dialed number (CLD/DNIS) from the call request with values in this field. If matched with a database entry, then the call will be authorized, rated and routed depending on the prefix specified in the <i>origin</i> field of this entry
origin	string	A target prefix used instead of the originally dialed number to authorize, rate and route the call. It can match a special recipient operator code maintained in tariffs (a symbolic name, e.g. Netcom) or a real destination from the recipient operator range. This prefix must exist in a routing tariff associated with this operator (vendor)
i_env	int	The unique ID of the virtual billing environment where the ported number belongs to. When a certain ported number must be checked only in a certain virtual billing environment it must have the <i>i_env</i> specified. The mapping table can contain globally applicable (used in all virtual billing environments) list of ported numbers – such entries will have the <i>i_env</i> field set to NULL. Note that in case of ambiguous match, the one with specified virtual billing environment is used

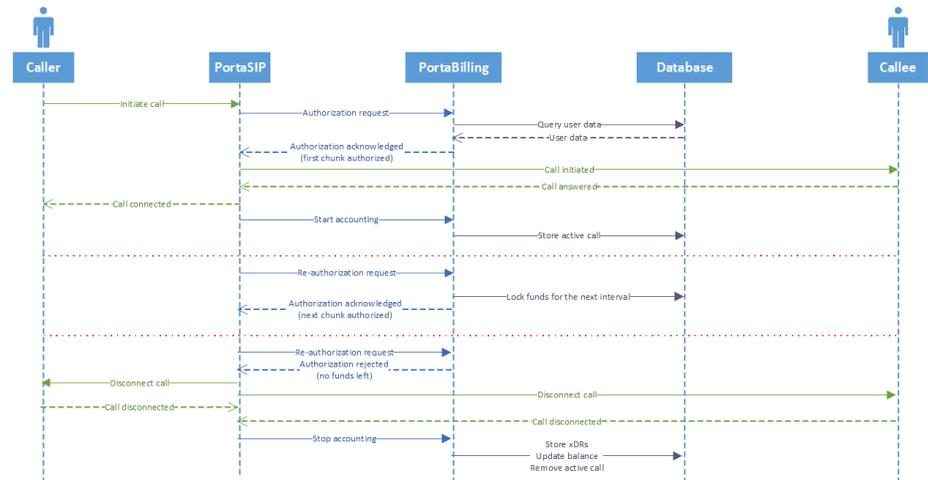
rate_match_code	string	By default the same target prefixes are used for routing and calculating your expenses for termination of calls to vendors and for billing customers. This might not be always convenient for customer tariff management because the rules for routing and billing vendors might be different from those used for billing customers. This field allows to specify prefixes for matching rates in customer tariffs. In this case prefixes from this field will be used for billing customers instead of the ones specified in the <i>origin</i> field.
rate_match_prio	string	Ported number priority (experimental.)

APPENDIX E. Billing diagrams

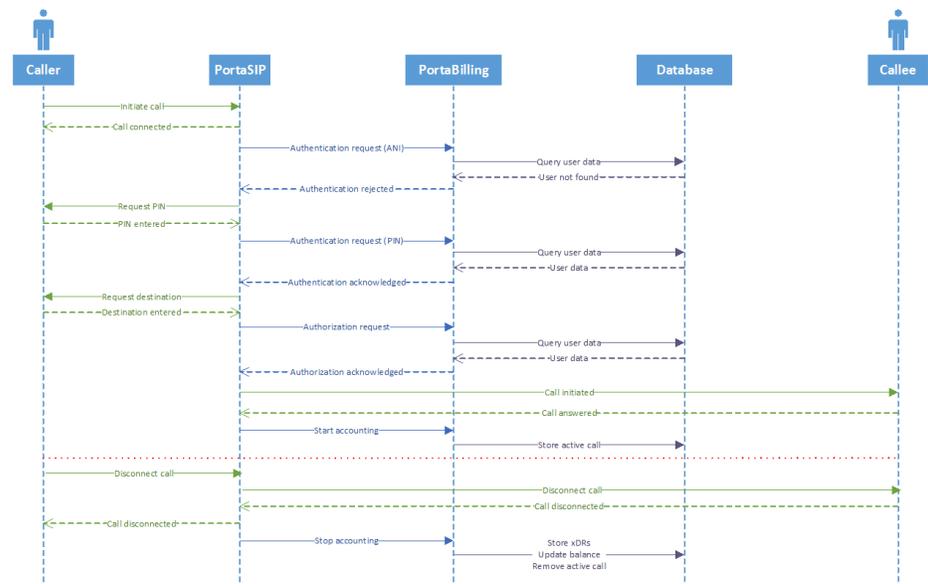
Call authorization and rating



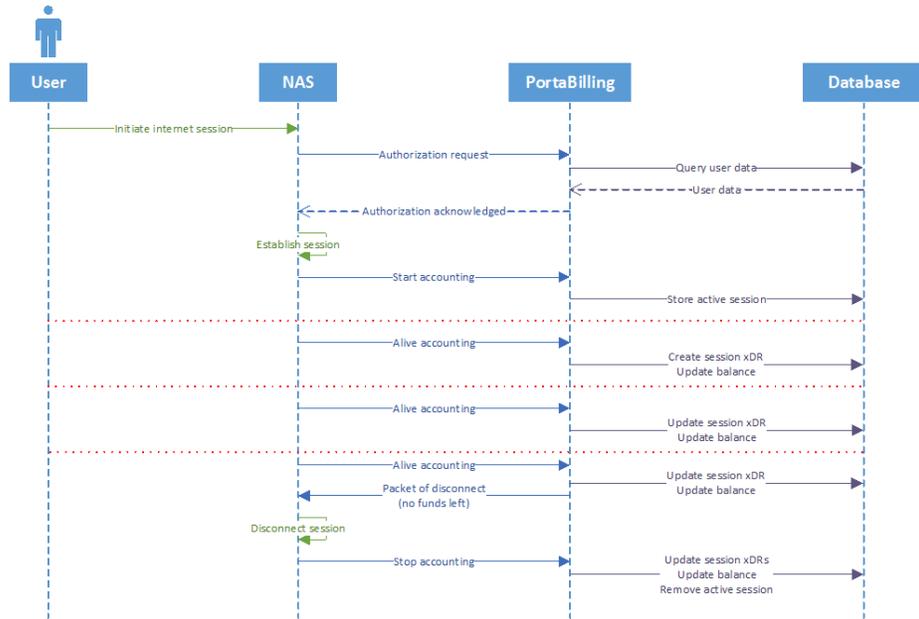
Call re-authorization and rating



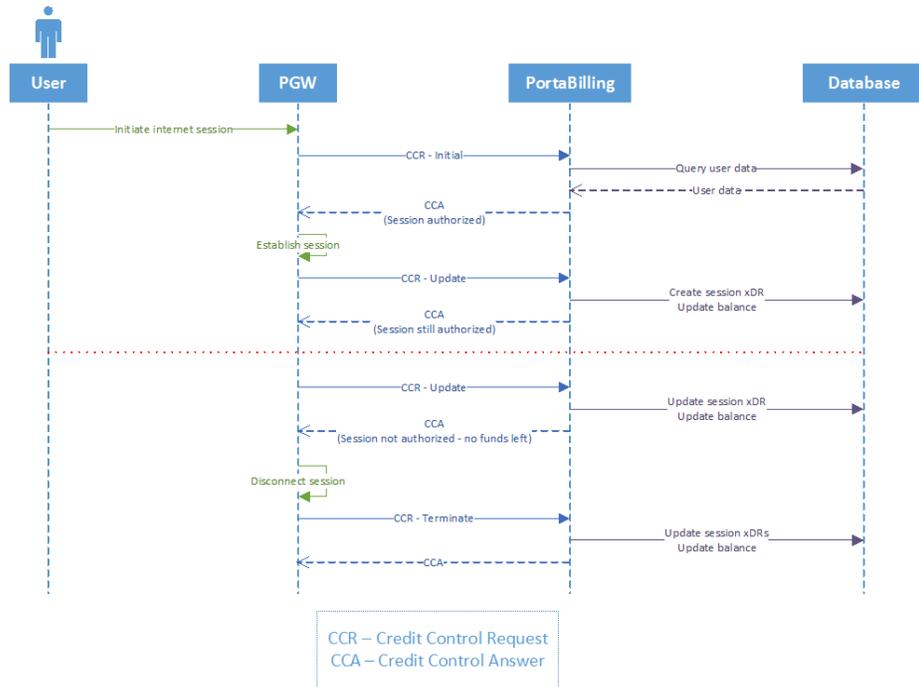
PINless dialing authorization and rating



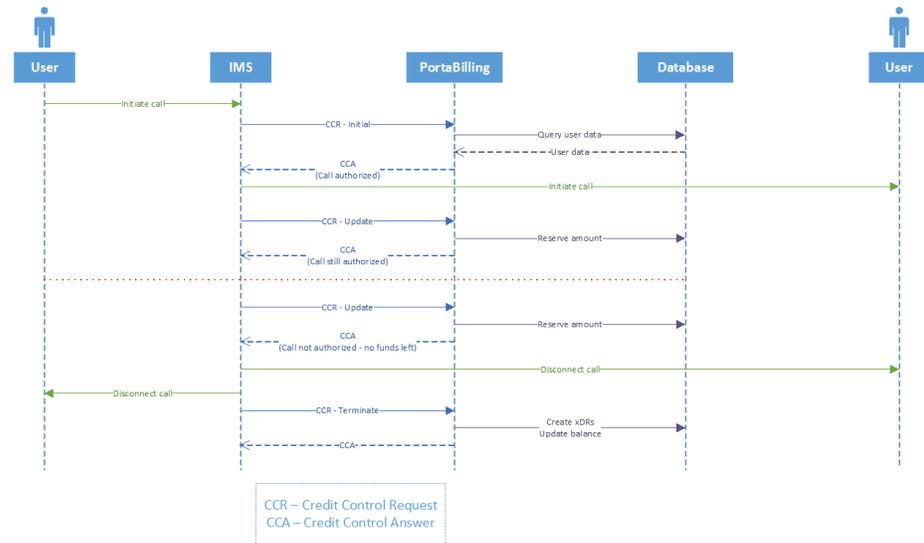
Broadband Internet access authorization and rating



LTE session



VoLTE session



APPENDIX F. Supported online payment processors

Name	Web Page	External	Token-based payments support
Assist	https://www.assist.ru		
AuthorizeNet	https://www.authorize.net		Yes
Bambora	https://www.bambora.com		
Chase Merchant Services (formerly Chase Paymentech)	https://merchantservices.chase.com		
Cielo	https://www.cielo.com.br		
CyberSource	https://www.cybersource.com/en-us/about.html		Yes
Docdata Payments	https://www.docdatapayments.com		
First Data Global Gateway Virtual Terminal (formerly LinkPoint)	https://www.firstdata.com		
Ingenico ePayments (formerly GlobalCollect)	https://www.ingenico.com/global-epayments		Yes

Moneris (formerly eSELECTplus)	https://www.moneris.com		Yes
Netaxept	http://www.paynet.no		
NetaxeptBBS	https://shop.nets.eu		
Nets (formerly Luottokunta)	https://www.nets.eu	Yes	
Payarena	http://www.unifiedpaymentnigeria.com	Yes	
PayGate (formerly VCS)	https://www.paygate.co.za	Yes	Yes
PaymentExpress	https://www.paymentexpress.com		Yes
PayNearMe	https://www.paynearme.com	Yes	
PayPal	https://www.paypal.com	Yes	
ProxyPay3	https://www.eurobank.gr		
Remita	https://www.remita.net	Yes	
SecureHosting	https://www.securehosting.com		Yes
SecurePay	https://www.securepay.com.au		
Skrill	https://www.skrill.com	Yes	
TransferTo	https://www.transfer-to.com		
Verisign (formerly PayflowPro)	http://www.verisign.com		Yes
WorldPay	https://www.worldpay.com	Yes	

NOTE: According to ChasePaymentech policy, you must pass the certification for your merchant account to use this payment processor.

APPENDIX G. Personal data that can be stored in PortaBilling®

The table below lists the kinds of personal data that can be stored in PortaBilling® and defines the availability for users with restricted access to personal data.

Entity	Personal information	Anonymized	Not available
Customer and reseller	Name	Y	
	Billing address information	Y	
	Address information including country, state, city, ZIP code	Y	

	Contact information such as company name, contact details, email, phone numbers, fax, BCC email	Y	
	Credit card details	Y	
	Additional fields in xDRs		Y
	Custom fields	Y	
	Sales agents' names		
	Self-care credentials	Y	
	Invoices	Y	
	Recipient's email for custom reports	Y	
	xDRs	Y	
Account	ID	Y	
	Service password	Y	
	Subscriber details such as name, email, phone numbers	Y	
	Address information	Y	
	Contact information	Y	
	Aliases	Y	
	Phone book details	Y	
	Follow-me lists	Y	
	Abbreviated dialing lists	Y	
	SIM card information: IMSI, MSISDN	Y	
	IP phone details: IP address and port, MAC address	Y	
	Location IP address specified for geo/IP verification	Y	
	xDRs	Y	
	Additional fields in xDRs		Y
	Call records		Y
	Voicemails		Y
Credit card details	Y		

Vendor	Self-care credentials	Y	
	Name	Y	
	Address information including country, state, city, ZIP code	Y	
	Contact information such as company name, email, phone numbers, fax, BCC email	Y	
	xDRs	Y	
	Self-care credentials	Y	
Representative	Name	Y	
	Address information including country, state, city, ZIP code	Y	
	Contact information such as email, phone numbers, fax, BCC email	Y	
	Initials	Y	
	Self-care credentials	Y	
	CC_staff	Name	Y
Address information including country, state, city, ZIP code		Y	
Contact information such as email, phone numbers, fax, BCC email		Y	
Self-care credentials		Y	
Administrative users (e.g. Helpdesk operators)		Name	Y
	Address information including country, state, city, ZIP code	Y	
	Contact information such as email, phone numbers, fax, BCC email	Y	

	Self-care credentials	Y	
Payment system	Merchant account credentials	Y	

APPENDIX H. Glossary / list of abbreviations

A

AAA – Authentication, Authorization and Accounting. A suite of network security services that provides the primary framework through which you can set up access control on your router or access server.

AD – Abbreviated Dialing

ADSL – Asymmetric Digital Subscriber Line

ANI – Automatic Number Identification

API – Application Programming Interface

ASP – Application Service Provider

ASR – Answer Seizure Ratio

B

BH – Busy Hour

BHCA – Busy Hour Call Attempts

BHCC – Busy Hour Call Completion

bps – Bits per Second

Bps – Bytes per Second

C

Call leg – A link or hop along the route from the calling party to the destination (called) party.

CDR – Call Detail Record

CLI – Command Line Interface

CLI, CLID – Calling Line Identification

CMS – Call Management Services

CPE – Customer Premise Equipment

CRM – Customer Relationship Management

CSP – Communications Service Provider

D

DCE – Data Communications Equipment

DHCP – Dynamic Host Configuration Protocol

DID – Direct Inward Dialing

DSL – Digital Subscriber Line

DTE – Data Terminal Equipment

E

E911 – Enhanced 911

F

FTTH – Fiber to the Home network

G

GK – Gatekeeper. The entity that maintains a registry of devices in an h323 network.

GPON – Gigabit Passive Optical Network

GSM – Global System for Mobile Communications

GUI – Graphical User Interface

GUID – Globally Unique Identifier (also called h323-conf-id). The identifier for end-to-end calls. A new GUID is assigned to each new call. There may be numerous calls within a single session.

H

HLR – Home Location Register

HTML – Hypertext Markup Language

HTTP – Hypertext Transfer Protocol

I

IMS – IP Multimedia Subsystem

iPaaS – Integration Platform as a Service

IP PBX – Internet Protocol Private Branch Exchange

IPTV – Internet Protocol Television

ISP – Internet Service Provider

ITSP – Internet Telephony Service Provider

IVR – Interactive Voice Response

L

LAN – Local Area Network

LCR – Least Cost Routing

LDAP – Lightweight Directory Access Protocol

LTE – Long Term Evolution is a standard for wireless communication of high-speed data for mobile phones and data terminals

M

MAC – Media Access Control

MAC address – The physical address of a device on Ethernet.

MMS – Multimedia Messaging Service

MNO – Mobile Network Operator

MVNO – Mobile Virtual Network Operator

N

NANP – North American Numbering Plan

NAS – Network Access Server

NGN – Next Generation Network

O

OEM – Original Equipment Manufacturer

OLT – Optical Line Terminal. The endpoint of the service provider for GPON service provisioning

ONT – Optical Network Terminal. The end user endpoint, installed on the customer premises

OS – Operating System

P

PBX – Private Branch Exchange

PDF – Portable Document Format

PIN – Personal Identification Number

POTS – Plain Old Telephony Service. Basic telephone service supplying standard single-line telephones, telephone lines, and access to PSTN.

PPP – Point-to-Point Protocol

PRI – Primary Rate Interface

PSTN – Public Switched Telephone Network

Q

QoS – Quality of Service

R

RADIUS – Remote Authentication Dial-In User Service

RAM – Random Access Memory

S

SBC – Session Border Controller

SG – Signaling Gateway

SIM Card – Subscriber Identify Module Card

SIP – Session Initiation Protocol

SMPP – Short Message Peer-to-Peer

SMS – Short Messaging Service

SOAP – Simple Object Access Protocol

SQL – Structured Query Language

SS7 – Signal System 7

T

Telco – A telephone company, such as LEC.

TFTP – Trivial File Transfer Protocol

U

UA – User Agent

URI – Uniform Resource Identifier

URL – Uniform Resource Locator

V

VoIP – Voice over Internet Protocol

W

WAN – Wide Area Network

X,Y,Z

xDR – Extensible Detail Record

XML – Extensible Markup Language