



PortaBilling



Administrator Guide



Maintenance
Release



Documentation

Copyright Notice & Disclaimers

Copyright © 2000–2015 PortaOne, Inc. All rights reserved

PortaBilling® Administrator Guide, October 2015
Maintenance Release 50
V1.50.09

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	6
Hardware and Software Requirements	7
Installation	8
1. System Architecture.....	9
Overview.....	10
Components	12
PortaBilling® Performance.....	15
PortaBilling® Web Interface.....	19
Web Cluster	20
2. System Concepts.....	22
Overview.....	23
Service Types and Services	24
Account and Account Types.....	27
Products and Service Bundling.....	32
Add-on Products	34
Retail Customers.....	42
Linked Customers.....	45
Distributors and Resellers	46
Vendors.....	58
Account ID-based Billing.....	60
Nodes.....	61
Virtual Environments.....	62
3. Rating and Invoicing.....	64
Destinations, Rates and Tariffs	65
Billing for Always-on Services	70
Session Billing Parameters	70
Call Billing Parameters.....	71
Volume Discount Plans.....	76
Bundle Promotions.....	87
Quantity Billing Parameters	91
Override Tariff	92
Overdraft Protection.....	94
xDR Recalculation	99
Subscriptions.....	101
Invoicing.....	114
Midterm Invoices.....	125
Out-of-Turn Invoice Generation.....	125
Processing Taxes.....	126
Manual Balance Adjustment.....	137
Outstanding Balance Adjustment.....	140
Payments.....	141
4. Voice Services	149
CDR Mediation	150
Call Routing	151

Number Translation.....	169
DID Provisioning and Billing.....	173
Local Number Portability and Other Special Cases.....	182
Favorite Number Billing.....	186
Registration for PINless Dialing Service via SMS.....	186
Bandwidth-aware Limitation of Allowed Concurrent Calls.....	187
Active Call Monitoring.....	188
IP Device Provisioning and Inventory.....	188
Using Different Protocols (SIP and H323).....	191
5. Converged Services.....	192
IPTV.....	193
Broadband Internet Access Service.....	194
LTE Service Provisioning and Billing.....	200
WiMAX Services.....	201
Integration with Bridgewater Service Controller.....	203
Mobile Virtual Network Operator.....	204
Integration with AudioCodes Mobile Application.....	205
Integration with Cisco PDSN.....	207
Integration with CISCO SCE Service Control Engine.....	208
Mobile Payment Remittance.....	209
6. Entering Data in PortaBilling®.....	214
Cost / Revenue.....	215
PortaBilling® Data Model Concept.....	215
Billing Entities Hierarchy.....	217
7. Management and Monitoring Tools.....	218
Rate Upload.....	219
Date and Time Information.....	220
Online Web Signup.....	221
Custom Reports.....	223
Trouble Ticketing System.....	227
PortaBilling Server Health Monitoring.....	229
VoIP Network Performance Statistics.....	231
Billing Statistics.....	234
Encrypted Passwords.....	239
8. How to	241
Log in to PortaBilling® After Installation.....	242
Terminate a Customer, Including All His Accounts.....	242
Locate the h323-conf-id for a Call.....	242
Troubleshoot an Incorrectly Billed Call.....	243
Make the Periodic Payments Tab Appear in the Customer / Account Info.....	245
Make the Custom Fields Tab Appear in the Customer / Account Info.....	245
Make a Custom Report from PortaBilling®.....	246
Use ODBC to Connect to PortaBilling®.....	246
Force PortaBilling® to Disconnect after a Customer Calls over his Credit Limit.....	252
Create Accounts to be Used for SIP Services.....	252

Integrate PB Logins to Your Website.....	252
Add Your Company Name and Logo to Every PortaBilling® Web Page	253
Add a Reseller’s Logo to the Reseller’s Customers’ self-care interfaces.....	257
Automatically Purge Old xDRs to Reduce the Size of the Database	258
9. Appendices.....	259
APPENDIX A. Supported RADIUS RFCs.....	260
APPENDIX B. Change-of-Authorization (CoA) and Disconnect Message (DM) Packets Supported by PortaBilling®	260
APPENDIX C. The Number_Portability Table Description.....	261
APPENDIX D. Glossary / List of Abbreviations.....	262
APPENDIX F. Hotlining Parameters	267

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. The last section of the document answers the most frequent questions users ask after running PortaBilling® for the first time.

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

- Five (5) Linux Servers for the PortaBilling®:
 - Real-time billing engine (RADIUS).
 - Main database server.
 - Web server (web interface for administrators and self-care portal, XML API services, periodic tasks, invoicing).
 - Replica database server.
 - 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® Switching Server and the PortaSIP® Media Server need to be processed), and 250 GB of free disk space if you are using the PortaBilling® standalone.

- A 64-bit processor (Xeon, Opteron). Additional processors are recommended for networks with a high call volume.
- 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 XP or above, Linux/BSD, Mac OS X 10.6 or above.
- **Web browser:**
 - Internet Explorer 8.0 or above, Mozilla Firefox 24 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 1024 × 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.

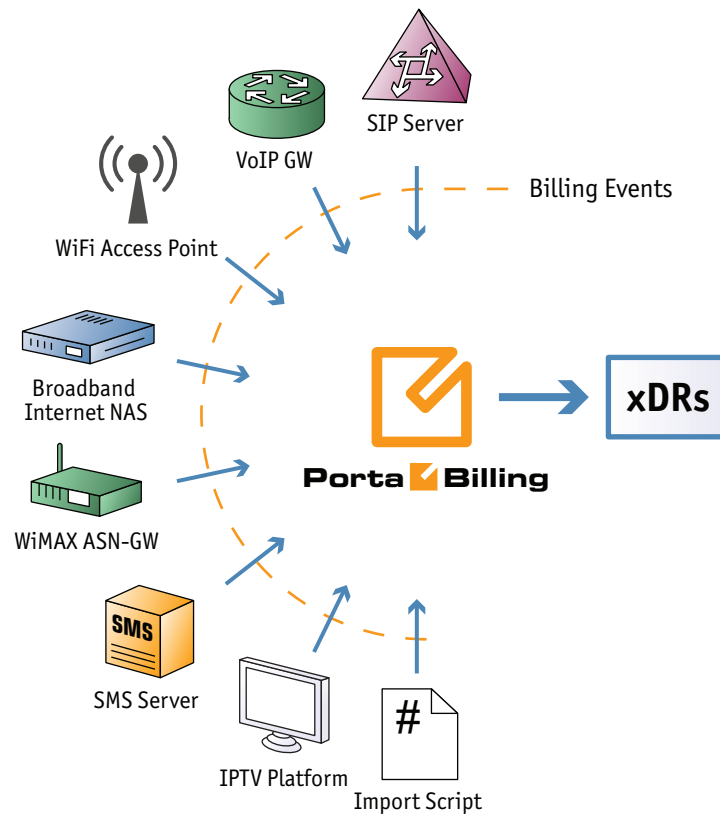
PortaSwitch® installation and configuration are automated and integrated within the main installation process. This allows you to install a completely functional PortaSwitch® environment from scratch in less than one hour!

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

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

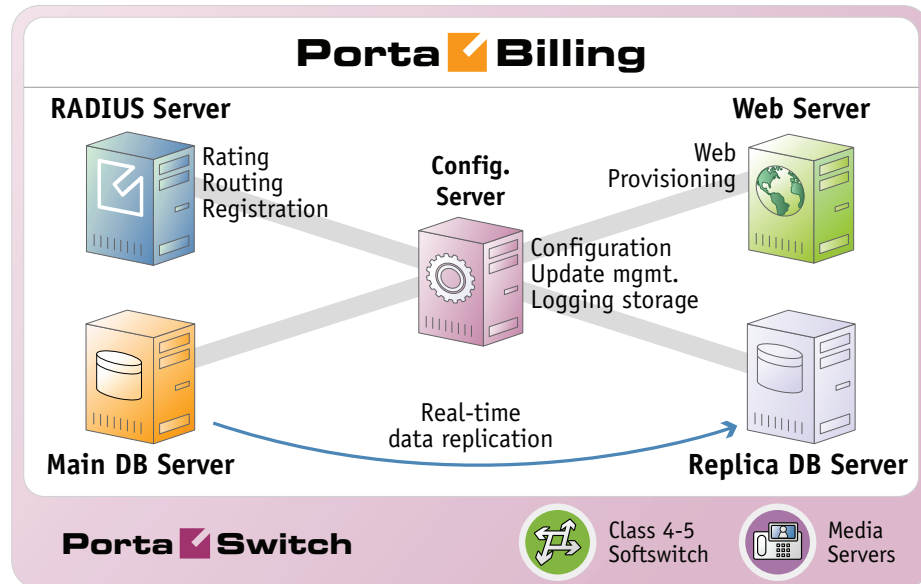
¹ For xDRs for telephony services, the previous term CDRs (Call Details Records) is often used.

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.




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 five main logical components: the main database server (lower left), the billing (RADIUS) server (upper left), the replica database server (bottom right), the admin/web server (upper right), and the configuration/log collector server (center). Typically, the billing, web and two database servers are installed on four separate physical computers so as to provide data redundancy and load sharing. Although a “shared” combination (where some servers, e.g. RADIUS 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.

RADIUS Server

The RADIUS server hosts the radius daemon, which has a billing engine embedded into it. So after it receives a request from the node it passes processing on to the billing engine and sends back a response.

Billing engine

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 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 RADIUS servers. 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® 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 is, of course, dependent on the clearly defined technical parameters of your car, such as engine horsepower or the car body’s total weight. So, in general, the more powerful car you buy, the faster you will be able to move. However, the actual speed of the car at any given moment depends on many other criteria: how many passengers are in the car, how much weight there is in the trunk, what the road conditions are (a paved highway vs. a gravel road), and so on.

The situation with billing systems is quite similar. The questions asked are: “How many customers can I have?”, “How many concurrent calls can it handle?”, or “How many messages or call minutes per month can I transport on my network?” The actual numbers depend not only on the PortaBilling® functionality and your hardware’s performance, but also on your billing configuration, traffic patterns, and the types of services provided. Let’s take a look at the details.

The only substantial performance criterion for PortaBilling® is the speed with which it can process requests from your network. PortaBilling® can process more than 300 RADIUS requests per second. What exactly does that mean? Every call attempt or completed call on your network generates several RADIUS requests, with the total number of requests and their type depending on the kind of service provided. For prepaid VoIP services, which involve call routing, this translates to about 50 call initiations per second, while for postpaid VoIP services there are about 75 call initiations per second. This means that 50–75 users can start a new phone call on your network each second. The exact figures will depend on the complexity of your billing configuration (e.g. how many rate plans, individual rates in a rate plan, or carriers you have). Since it cannot be known beforehand what combination of different services you will be

providing on your network, it only makes sense to refer to a conservative **estimate** of the PortaBilling® processing speed, which has been calculated at **50 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 aggregated call processing speed of 50 call attempts per second, an average call duration of 5 minutes, and a call success rate of 50% (the industry norms), it means that 50% of the 50 call attempts per second will succeed. So 25 calls will be connected, while the same amount of previously connected calls will be disconnected. Since the average call duration is 300 seconds, this means that at all times approximately $25 * 300 = 7,500$ calls will be in a “connected” state. If your ASR changes, it will have an immediate impact on the number of concurrent calls, e.g.:

- ASR 25%: $(50 * 25\%) * 300 = 12.5 * 300 = 3,750$ concurrent calls.
- ASR 60%: $(50 * 60\%) * 300 = 30 * 300 = 9,000$ 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 differing amount of calls in peak and off-peak hours) and your call success rate. We can estimate this based on typical traffic patterns (as indicated below) with 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 hours, peak hours, and extraordinary traffic splashes (e.g. New Year’s Eve). Thus the total performance capacity must be greater than the maximum anticipated BHCA (busy hours call attempts). On average, extreme traffic peaks (such as Christmas, New Year’s, etc.) will have 2 to 3 times more call attempts than your normal peak hours, so approximately $50 / 2.5 = 20$ call attempts per second can be expected during your peak hours. Using the same “industry norm” figures as above (50% ASR, 5 mins ALOC), PortaBilling® can process $(20 * 50\%) * 5 * 3600 = 180,000$ minutes during your peak hours. We can estimate that your total amount of daily traffic will be about 6 times the amount of traffic during one peak hour, i.e. $180,000 * 6 = 1,080,000$ minutes per day. Then, multiplying the daily amount by the number of “business” days in a month, we get $1,080,000 * 20 = 21,600,000$ minutes. This is why (to be on the safe side and allow for some reserve capacity) PortaBilling® is poised to handle monthly traffic of over **20 million minutes**.

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.

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® Oracularius and PortaSwitch® Procinctus products.

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

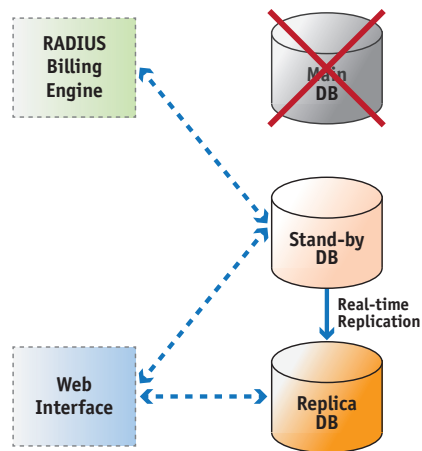
Since the load is distributed among all servers in the cluster, the total performance capacity grows linearly with adding more servers into the cluster. So if we add to a PortaBilling®, capable of **50** call attempts per second, another identical server and connect these two into a cluster – the performance of this two node cluster will be about **100** call attempts per second and by adding a third node it can be increased to **150** 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 RADIUS 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.

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.

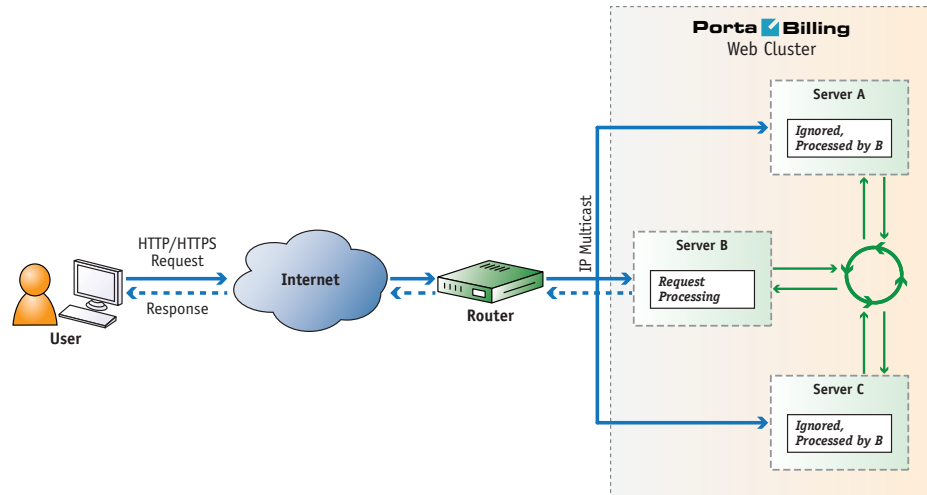


For improved security, the PortaBilling® web interface is divided into several separate areas, each serving its own purpose. Thus a login name and password which allow a user to access the customer self-care interface cannot be used to log in to the admin interface; the login name and password for the account’s (end-user’s) self-care interface are not applicable to customer self-care interface; and so forth.

Web Cluster

The purpose of Web Cluster is to allow 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 between multiple servers, allowing the system to easily scale up to handle a higher total number of users who are accessing the system via web interface or external applications, or accessing the data via XML API.

For outside users (or applications) the cluster is accessible via a single IP address. For instance, if the cluster 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 the user types “www.mycall.com”, his request will be 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 on external applications or via the end-user’s web browser.



All incoming sessions are randomly distributed among available nodes based on the source IP address and port using IP multicast (RFC 2365). This means that a single HTTP request (although it may be transmitted in multiple IP packets) will always be processed by the same node in the cluster. If one node in the cluster is down, the remaining nodes will perform the cluster reconfiguration and share the incoming traffic.

Delivery of the incoming packets to all servers in the cluster requires that all servers are connected to the same LAN segment and that the router, connecting that network to the Internet supports the IP multicast feature. The cluster starts with a minimum number of two servers. The number of Web servers that can be added to the cluster is unlimited.

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 be equal to or greater than the measurement unit, and the billing unit must be equal to or greater 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.

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.
Dialup Internet	S	Dialup 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.
Messaging Service	Q	Rating messages (text, SMS, MMS, other); charges are based on the number of messages sent.
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.


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.

This 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.

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). In this case, incoming calls to the alias will be routed to the main account. Outgoing calls using the alias cannot be made.

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. You can block or unblock, change the balance, or perform other operations for either the whole batch or a portion of the batch. 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.

Account Life Cycle

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

In addition, there are two parameters that define the life cycle of an account:

- Every account has an **activation date**. This defines the date of the 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).
- **Expiration date** defines the date after which the account can no longer be used. There are three options for defining it:
 - **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.

These options can be chosen together or separately. They have equal priority: the account expires according to which event occurs first. The expiration date is optional, so you may also create accounts with no expiration date.

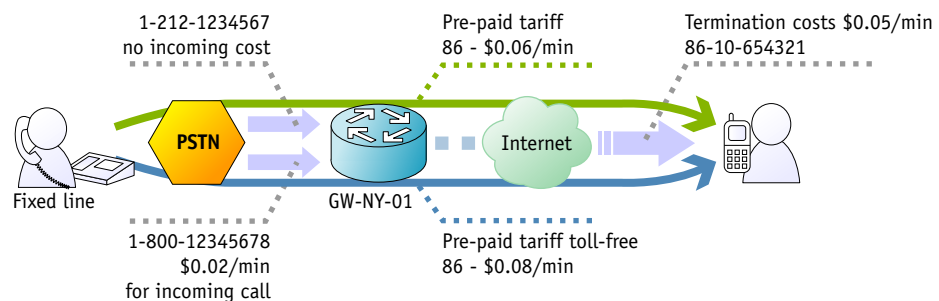
Another important point regarding expiration options is that a voucher recharge operation extends the life cycle of an account. Thus, if an account is adjusted to be available for 60 days after first use, and the account's owner recharges the account 57 days after the date of its first use, the service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

Products and Service Bundling

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** tab) 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, where the call originates on the PSTN network). 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.

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 WiMAX network and be charged for time spent online, and make calls from their mobile phone (or land-line) via PINless dialing service (dialing an access number where the system will automatically recognize them by the number they are calling from, and then entering the destination number). You may assign a different rate plan for each service, as shown; for PINless dialing services there is even a different rate plan for calls made from fixed phone line and from mobile phone.

The screenshot shows the 'Edit Product "MegaStar Product"' interface. At the top, there are navigation buttons: Save, Save & Close, Close, Rate Lookup, and Clone. On the right, there are links for Logout and Log. The main form contains fields for Product Name (MegaStar Product), Currency (USD - US Dollar), Product Name visible to End User (MegaStar Product), and Managed By (Administrator Only). Below these are radio buttons for Product Type: Main Product (selected) and Add-on Product.

The interface has several tabs: Included Services, Service configuration, Usage Charges, Volume Discount, Recurring Charges, Additional info, and Notepad. The 'Usage Charges' tab is active, displaying a table with columns for Edit, Log, Service Used At (Node, Access Code, Service), Charges Applied (Tariff, Advanced Config), and Delete.

		Service Used At			Charges Applied		
Edit	Log	Node	Access Code	Service	Tariff	Advanced Config	Delete
Service type: Voice Calls							
		PortaSip	INCOMING	Voice Calls	MegaStar - Incoming		
		PortaSip	OUTGOING	Voice Calls	MegaStar - Outgoing		
		PortaUM	ANI#18661234555	Voice Calls	Prepaid Calling Cards		
		PortaUM	FIXED#18661234555	Voice Calls	Prepaid Calling Cards FIXED		
		PortaUM	MOBILE#18661234...	Voice Calls	Prepaid Calling Cards MOBILE		
Service type: WI-FI							
		Mikrotic.WiFi	Any code	Wi-Fi	WIMAX		

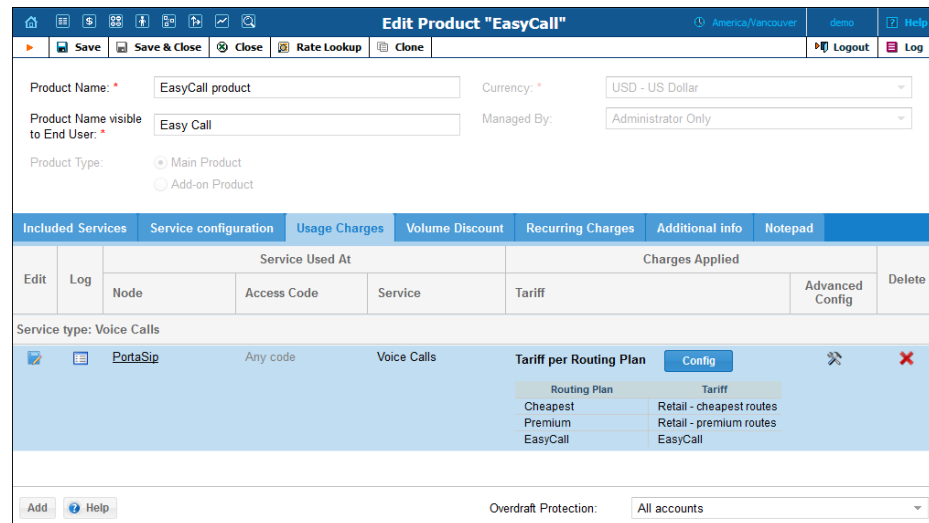
At the bottom, there are 'Add' and 'Help' buttons, and an 'Overdraft Protection' dropdown menu set to 'Debit accounts only'.

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** select menu.

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; and
- 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 are 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); and
- 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:

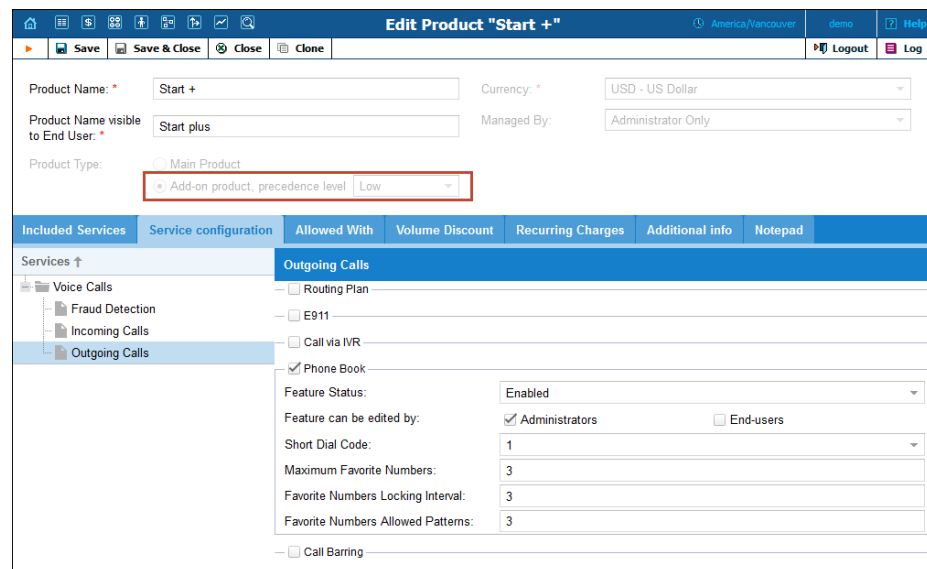
- ensures that customers are supplied with the quality services that they need; and
- decreases the workload on administrators as it is easier to manage those products in future.

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 **Service Features** tab 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:

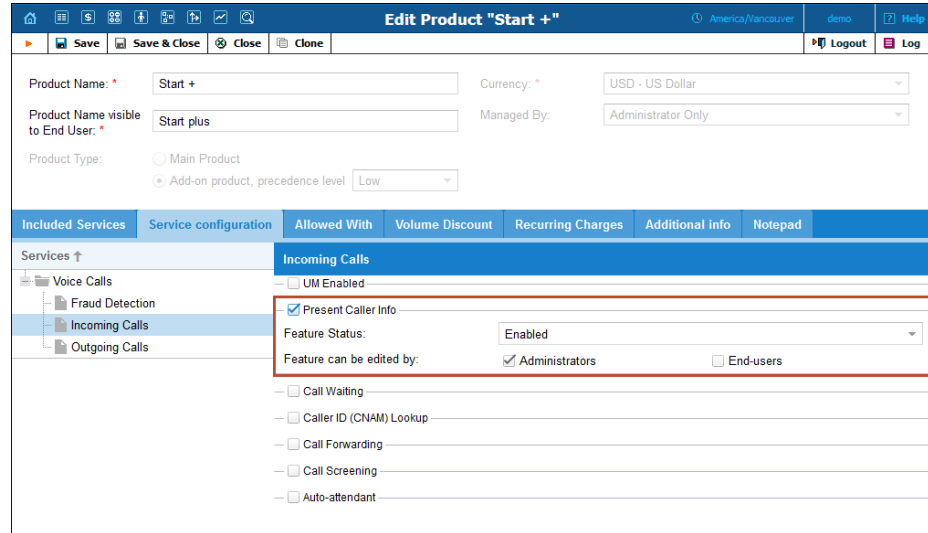
- **Precedence.** Add-on products have precedence over main products, so no matter what is defined within a main product it will be overridden by settings defined within the add-on products (this applies only to the options that are supplied with add-on products). In order to differentiate between the 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.



The screenshot shows the 'Edit Product' interface for a product named 'Start +'. The product is configured as an 'Add-on product' with a 'precedence level' of 'Low'. The 'Service configuration' tab is active, showing a tree view of services under 'Outgoing Calls'. The 'Phone Book' feature is checked and enabled, with a status of 'Enabled'. The 'Feature Status' is set to 'Enabled', and the 'Feature can be edited by' is set to 'Administrators'. The 'Short Dial Code' is '1', 'Maximum Favorite Numbers' is '3', 'Favorite Numbers Locking Interval' is '3', and 'Favorite Numbers Allowed Patterns' is '3'. Other features like 'Routing Plan', 'E911', 'Call via IVR', and 'Call Barring' are unchecked.

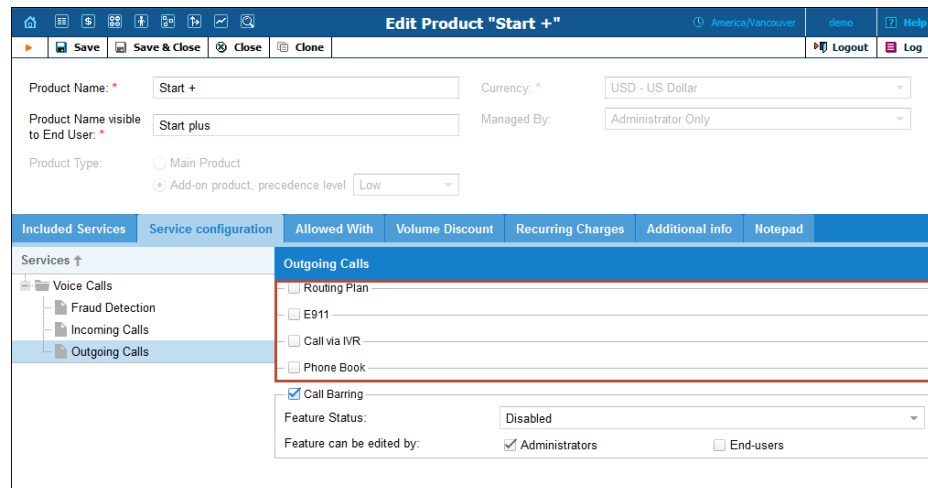
For example, an account with a “Basic Call” main product is assigned a “Smart Call” (low precedence level) add-on product with the **Conferencing** feature disabled. Then the customer decides that he needs the conferencing service, so he orders the “Conferencing+” (high precedence level) add-on product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



The screenshot shows the 'Edit Product "Start +" interface. The 'Included Services' tab is active, and the 'Incoming Calls' service is selected. A red box highlights the 'Present Caller Info' feature, which is checked. Below it, the 'Feature Status' is set to 'Enabled' and 'Feature can be edited by' includes 'Administrators'.

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).



The screenshot shows the 'Edit Product "Start +" interface. The 'Included Services' tab is active, and the 'Outgoing Calls' service is selected. A red box highlights the 'Call Barring' feature, which is checked. Below it, the 'Feature Status' is set to 'Disabled' and 'Feature can be edited by' includes 'Administrators'.

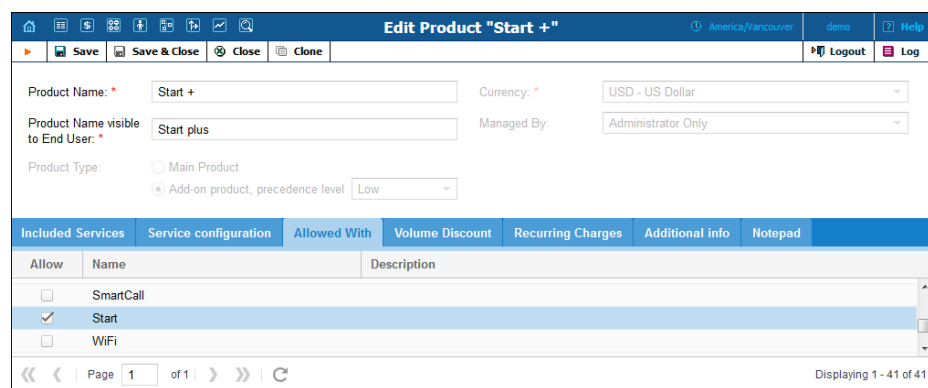
- **Service Attributes.** All attributes can be sorted into three categories:
 - Those that are defined on the product level only – they are usually the ones that have some business value and should be controlled by a product.
 - Those that are defined for every account separately, e.g. the list of follow-me 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 precedence level cannot occur. The system doesn't 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 precedence level, so saving this configuration will not be allowed since it is unclear whether the **Call Waiting** functionality should be enabled or disabled.

Add-on Products' Subscriptions

A feature and a subscription are attached to add-on products, so an administrator only needs to assign the add-on product to an account. The feature will be activated and a customer will automatically be charged for it.

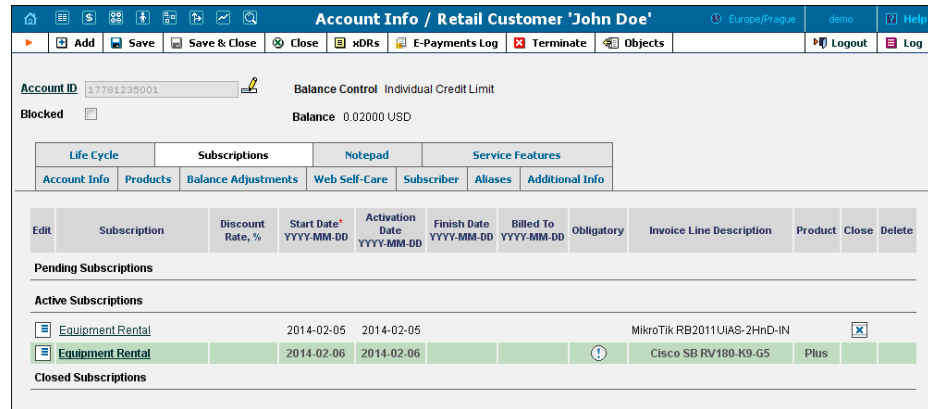
For example, a customer is assigned the “Call Waiting” add-on product with the Call Waiting feature enabled and a subscription associated with it that applies recurring charges for its usage. If the customer no longer needs Call Waiting functionality, the administrator can remove the add-on product from the account properties with a single mouse click.

Multiple subscriptions

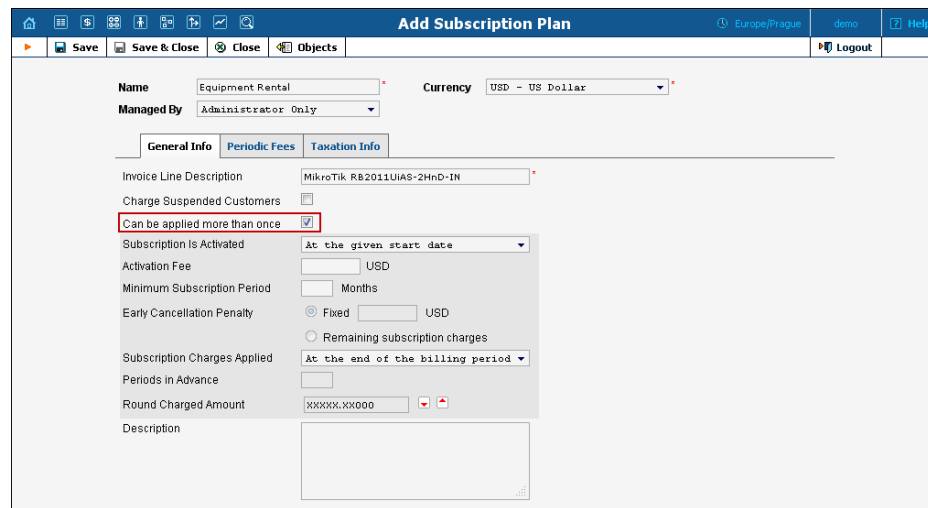
In most cases, every single subscription is charged only once even if this subscription is included in with several add-on products assigned to the account.

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. Another day he wants to rent a second router. Since this customer's account already has the subscription for renting a router, a new add-on product needs to be created with the same subscription assigned for the rent of the second router. The subscription has “Apply more than once” option enabled. According to these circumstances, the customer will be charged the same subscription twice: for the (first) already rented router and for newly rented one.



For this reason, there is an option that makes it possible to assign the same subscription multiple times to one account on the subscription plan properties.



This model of applying subscriptions will make it possible to:

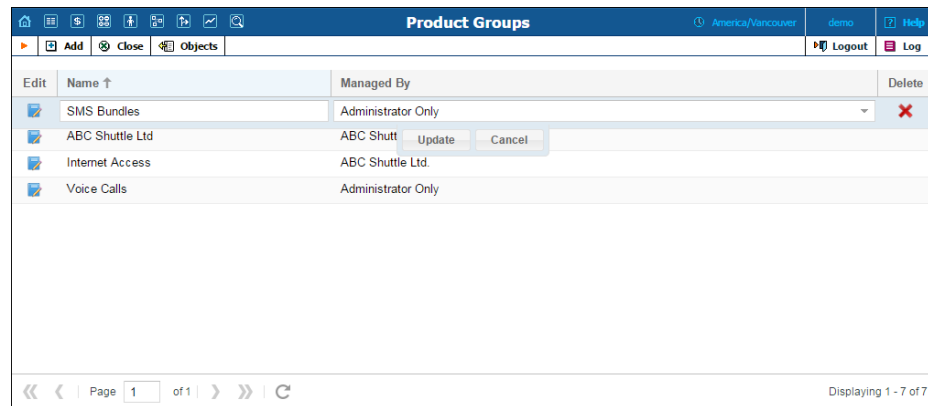
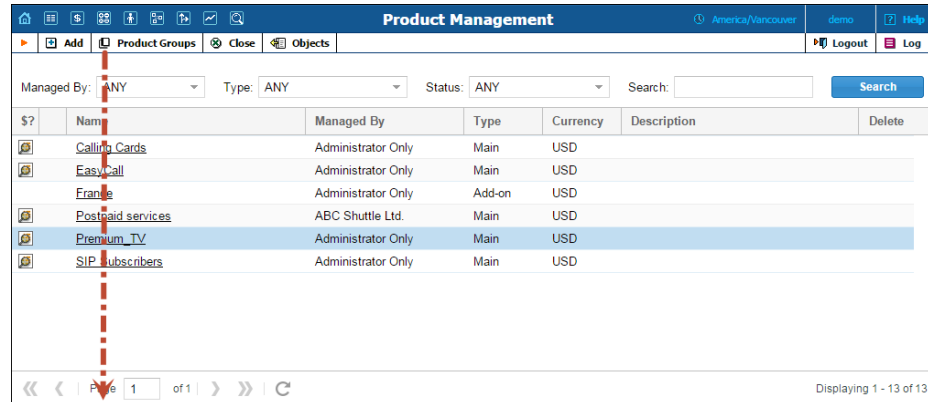
- provide revenue assurance, because using an add-on product guarantees that the feature will be activated, and the customer automatically charged for it; and
- save time, because everything can be done in one place using an add-on product.

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 helpdesk.

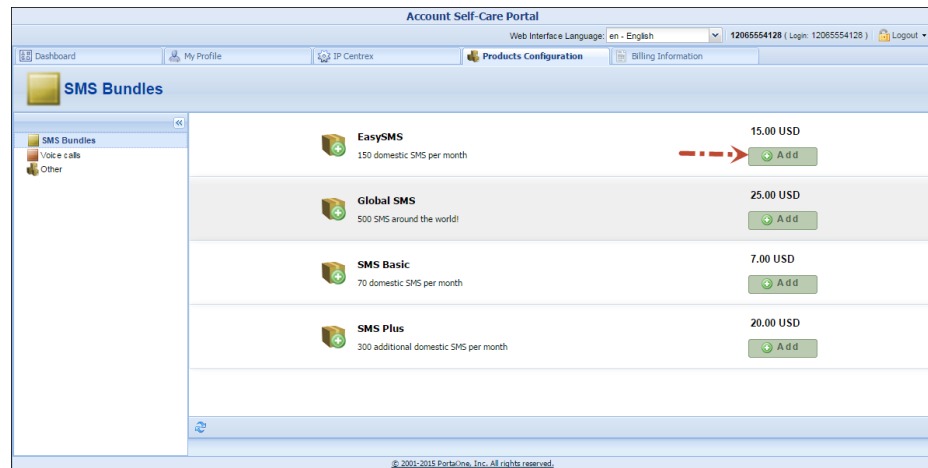


Then these groups are available to end users on the **Products Configuration** tab of the account self-care interface. A user can easily find a desired add-on by entering the corresponding group that the add-on belongs to.

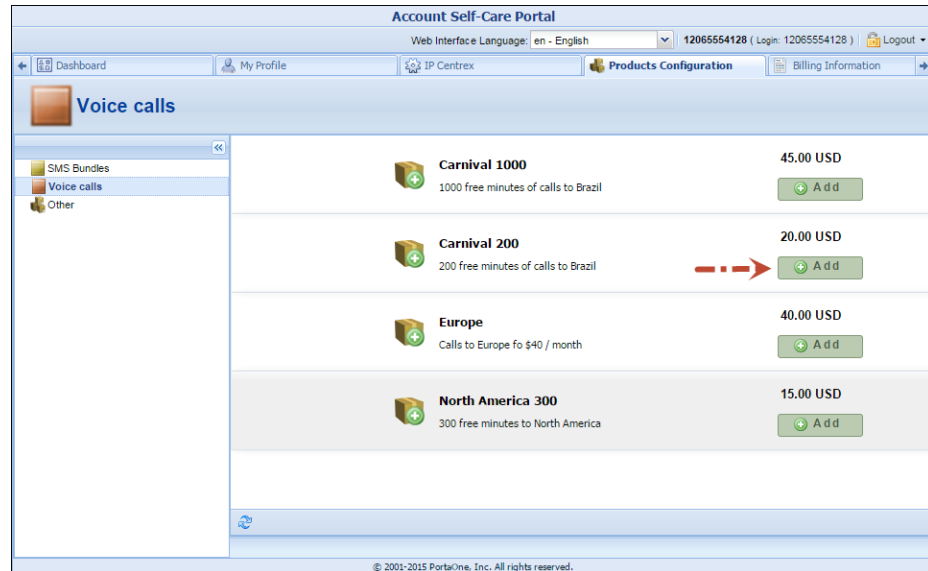
Consider the following example:

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.

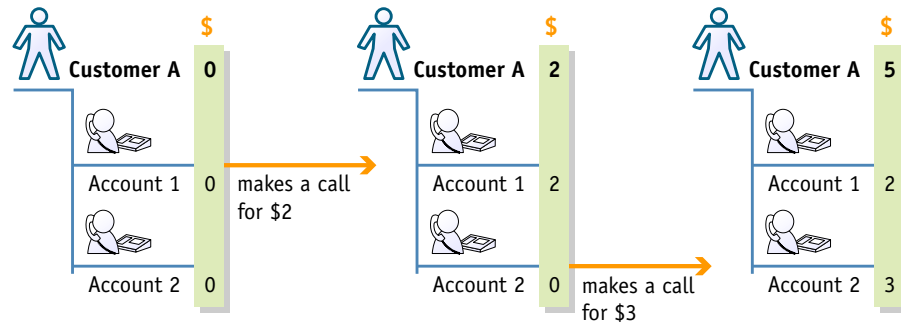
Thus, product groups enable end users to perform a quick and easy product search and self signup, thereby decreasing the load for your administrators and customer care staff.

Retail Customers

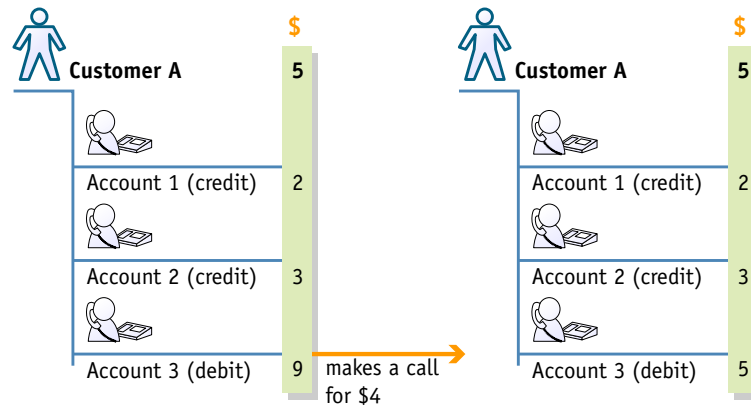
A customer is an entity (an individual or a company) who is responsible for using the service via one or more accounts. For instance, when a family signs up for residential VoIP services, a man, his wife and his brother will all have individual phone lines, each represented by an

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. This allows you to mix prepaid and postpaid accounts under the same customer – e.g. your postpaid customer who has multiple IP phones can purchase additional prepaid cards for his temporary workers.



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	Postpaid	Credit subordinate *	IP Centrex, messaging service, triple-play service provisioning for residential customers

all users under that customer.			
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.

Linked Customers

With the *linked customers* functionality, 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 one linked customers' group:

- **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.

The *linked customers*' functionality 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.

Distributors and Resellers

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.

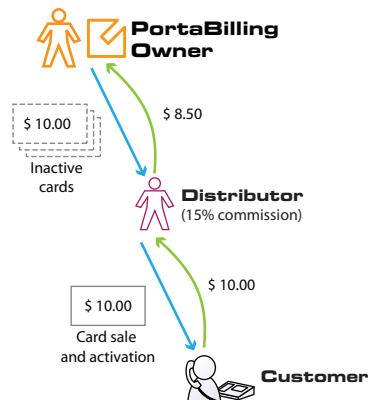
Customers can purchase new products, settle their invoices or refill their prepaid accounts by paying cash to a distributor (for example, the owner of the grocery shop opposite your house could be a distributor). After the distributor collects the money, he delivers it to the ITSP, minus his commission. 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. He either delivers to the end user a tangible product (e.g. calling card) supplied to him beforehand in bulk, or, in the case of an intangible product (e.g. registering a customer's cell phone number so he can make

PINless dial calls), he 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.

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.

Customer ID	Customer's name or company	Address	Accounts	Balance, USD	Amount Due, USD	Date	Status
John Doe		1410 ****, New York, NY 11229	****7885 ****9754	24.73	24.73		

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 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.

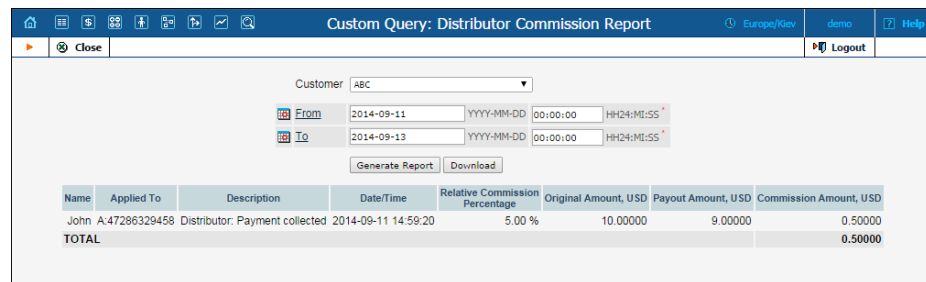
When enabled on the Configuration Server, commission reports are accessible for distributors via the self-care interface.

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.

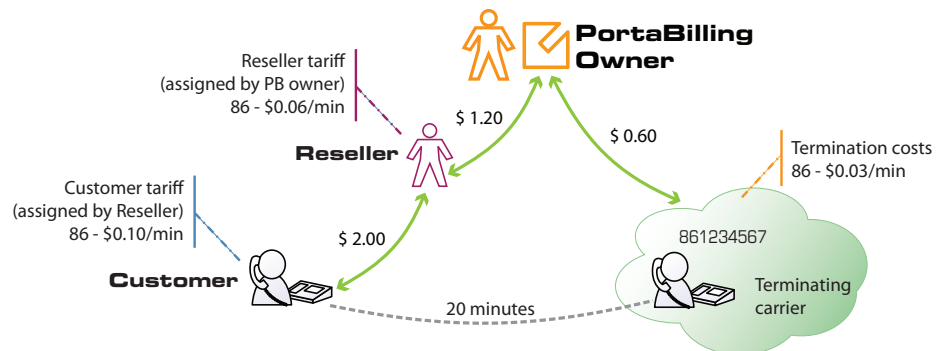
Thus, using commission report tools, distributors have a clear idea of the profits received from performing business operations.

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

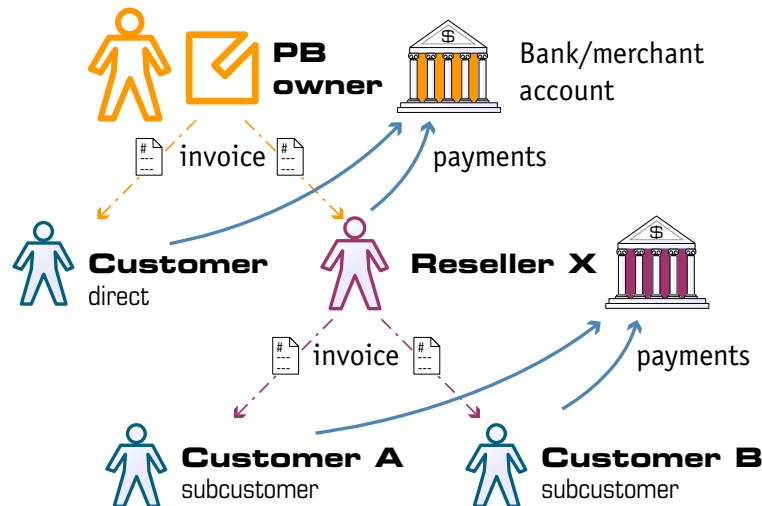
Resellers

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 USD/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 USD/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

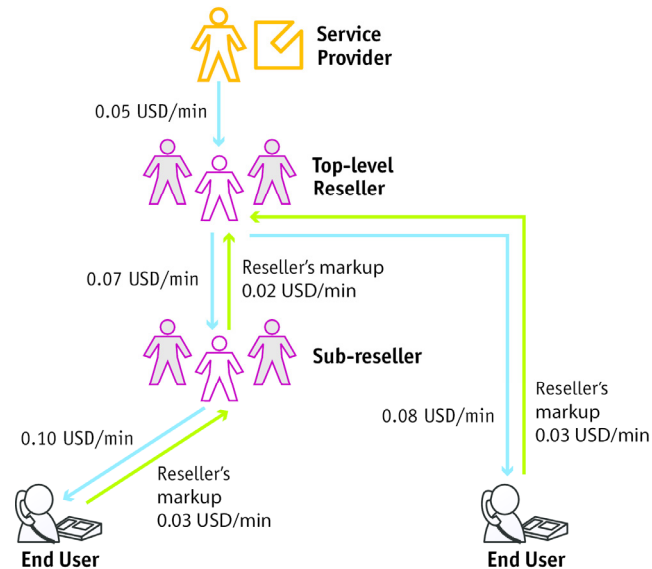
To reach more customers in the market many ITSPs sell their services through a multi-level network of resellers. To implement this functionality in PortaBilling®, the “carrier – reseller – end user” model has been extended to a multi-level reseller structure where the entire chain of resellers is virtually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

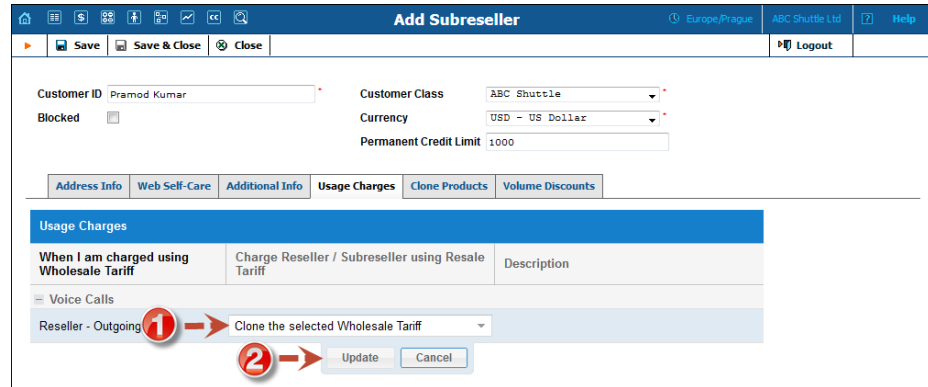


Managing tariffs

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.

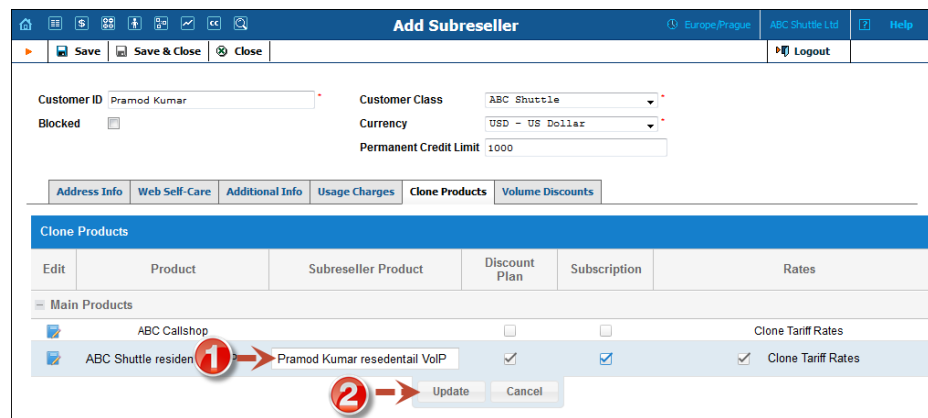
Wholesale tariffs are assigned to a reseller during their creation. *Resale* tariffs are configured as follows: when a reseller creates his subreseller, he clones the existing *wholesale* tariffs by applying a profit markup (a percentage) to them. Mapping between reseller's wholesale and resale tariffs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



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.

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 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.

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.

Customer status

A customer may have the following status:

Status	Description
 Credit exceeded	This indicates that the customer’s balance is above his credit limit, so he will not be able to make any outgoing calls, unless these are free calls (e.g. toll-free numbers).
 Payment frozen	This indicates that the customer’s auto-payment has been suspended due to repeated errors (e.g. credit card cancellation).
 Blocked	The customer has been blocked by the administrator; no call services will be provided until the administrator removes the block. The customer still has access to his self-care pages. Blocked customers are not subject to maintenance charges, but subscription fees still apply.
 Suspended	Services for this customer have been suspended because of an overdue invoice. Once payment is received (whether the customer makes an online payment, or a periodical payment is applied, or the administrator enters a cash/check payment manually) in an amount covering the overdue invoice, the suspension is automatically lifted.
 Closed	The customer’s account has been closed, and is kept in the database only for informational/regulatory purposes. No further operations are possible with this entity.

In the tables below you can see whether services are allowed or denied depending on customer status and type of service. The **Overdraft Protection** setting for the service at the product level is the following:

- The **Check Minimum Funds** option is set to *Do not check*.

Customer Status	Type of service	
	Toll-free	Chargeable
Credit exceeded / Zero Balance	allowed	denied
Payment frozen	allowed	allowed
Blocked	denied	denied

Suspended	denied	denied
Closed	denied	denied

- The **Check Minimum Funds** option is set to *More than 0*.

Customer Status	Type of service	
	Toll-free	Chargeable
Credit exceeded / Zero Balance	denied	denied
Payment frozen	allowed	allowed
Blocked	denied	denied
Suspended	denied	denied
Closed	denied	denied

Customer Classes

Customer classes allow you to define 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).

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.



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.

Account ID-based Billing

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.

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

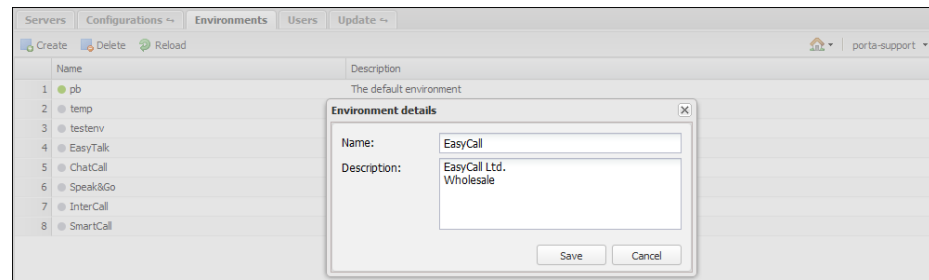
It is possible to have multiple isolated PortaBilling® instances, or virtual environments, within the same PortaBilling® installation. This functionality allows you to use PortaBilling® as an ASP (Application Service Provider) platform, or to create test environments which do not affect billing production 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 which you will use to login and then create other users, tariffs, customers, and so on. By default, the username of this user is `<env>-root` and the password is the same, where `<env>` is the name of the environment you have just created. So if you create the environment **easycall**, to log in to it you will use the username `easycall-root` and the password `easycall-root`.

3 ■ Rating and Invoicing

Destinations, Rates and Tariffs

Destinations, rates and tariffs are the essential parameters which define how a certain event should be billed. It is very important to find a correct set of billing parameters and rules (rates) for every event, and this is done by matching a service identifier with the available rates in the rate plan (tariff). The service identifier specifies exactly how the service was 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 likewise be a destination phone number, while for services such as a pay-per-view movie it could be the movie category (“New Release”, “International” or “Classic”). Since the rating of telephony numbers is the most complex case, it 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 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 have to be entered into the system before you can use these prefixes. This prevents errors and helps you to improve data quality. It is recommended that all of your prefixes be defined in the E.164 format.

It is virtually impossible to have an “ultimate” destination list that would contain all prefixes for the entire world. First of all, it is quite difficult to gather and maintain such information. Second, and even more important, is the fact that having “all possible” prefixes would not offer us 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 will need to 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, as it will not ever actually be used.

Therefore, we recommend that you follow a “required minimum” approach: only create those destinations for which you will 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.

Destination group set "MCI"



Destination group set "Retail"



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

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 USD/min during peak time and 0.07 USD/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.

Billing for 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).

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, as explained in the following section.

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® Web Reference Guide](#) manual 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 `Multiple OffPeak Periods` 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

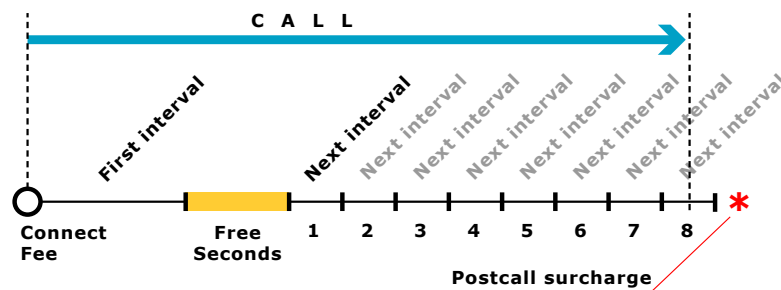


Figure 3-1

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 in call rating. You create a call formula which is then applied 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** – Charged part of remaining call duration; the parameters for this element are:
 - **Duration** – Rounding period (in seconds).
 - **Count** – Number of rounding periods in the interval.
 - **Price per minute** – Automatically prorated according to the rounding period duration.

The charge for the interval is calculated as:

$$\text{Total} = N * (\text{Duration}/60) * \text{Price_per_minute},$$

where *N* is either equal to *Count* (if *Call_Duration* is more than *Count * Duration*) or else *Call_Duration/Duration* (rounded up).

Thus you can create intervals such as “First 10 minutes – price per minute is \$0.05 and we will round up to 30 seconds” – in this case, the duration will be 30, the price per minute 0.05 and the count 20 (because there are twenty 30-second intervals in 10 minutes).

- **Fixed Surcharge** – Fixed amount to be added to the total call cost.
- **Relative Surcharge** – The total call cost (at this moment) will be increased by *X*%.

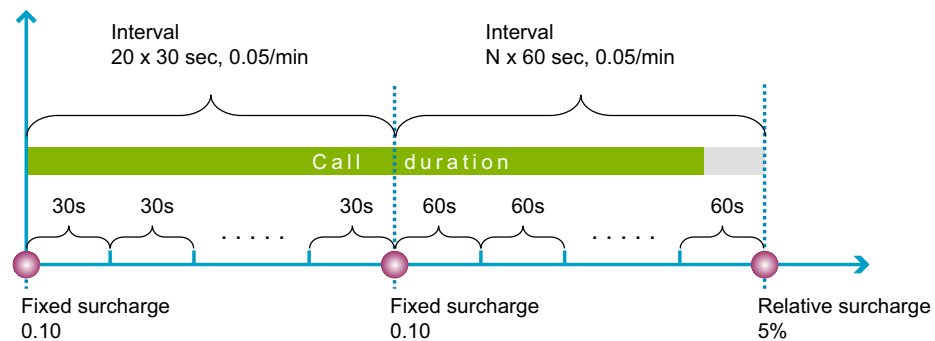
Formula elements are applied until there is no more non-charged call duration left. The surcharges which follow a certain interval are applied only if this interval has been fulfilled. “Fulfilled” means that the interval is covered completely, i.e. the remaining call duration was enough to cover the $\text{Count} * \text{Duration}$ 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 $N * 60$ 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 **Service&Rating** 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.

Volume Discount Plans

Sometimes you want to adjust rates based on the quantity of services your customer has already used. Volume-based discounts are ideal for this. The key features of volume-based discounts include the following:

- A discount is defined per destination group. For example, if you wish to provide special rates for calls to “ex-USSR countries” you can create a destination group that will contain prefixes for Russia, Ukraine, Belarus, etc. and define a single discount rule for this destination group.
- Multiple discount steps per destination group. For example, you can define a discount rule stating: the first 200 minutes free (100% discount), then up to 500 minutes at the normal rate (0% discount), then a 10% discount for every minute up to 1,000 minutes, and a 20% discount for every minute over 1,000.
- Discounts can be applied regardless of the time when the service is used. Alternatively, a different set of thresholds can be defined for peak and off-peak periods. It is also possible to define a volume-based discount which would only function during peak periods and not during off-peak ones or vice versa.
- 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).
- Several discount rules may be grouped into a volume discount plan. 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 (e.g. for calls to US&Canada or India) is represented by an individual discount rule. Taken all together, they will be grouped as discount plan “EasyCall,” which can then be applied to a specific account.
- For better management, you can associate a volume discount plan with a certain product. This discount plan will then be applied by default to every account using this product.
- You can assign several volume discount plans to an account with the help of add-on products. When several add-on products with volume discount plans are assigned to an account, the account receives a combination of those discount plans.
- You may also assign an individual volume discount plan for a particular account. This specifically assigned volume discount plan defines how discounts defined within other volume discount plans are combined and applied to an account.
- The usage counters that define which discount level is available, can periodically reset (allowing you to provide services such as “200 free minutes per month”) or can simply work as a one-time promotion (e.g. “Sign up today and receive 1000 minutes of discounted calling.”).

Discount calculations

The discount is applied to every minute which **exceeds** the threshold. Let's look at the following example:

- You have the discount rule “first 200 minutes – normal 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).

Usage periods

In order to provide more flexibility for creating new products, a discount has its own “usage period” for each destination group – the period that 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.

Applying discounts to customers or accounts

Discount plans may also be assigned at the customer level. In this case, they will apply to calls made by any of the customer's accounts.

If both the account and the customer have discount plans assigned for the **same** destination group, the discount plans are sorted by their precedence. The precedence is determined as follows (from the highest to the lowest):

- Discount plans assigned directly to an account;

- Discount plans assigned to add-on products. If more than one add-on product is assigned to an account, they will be sorted according to their specified priority;
- Discount plans assigned to the main product;
- Discount plans assigned to a customer.

The volume discount plan with the higher precedence in the list defines the way the next applicable discount plan with a lower precedence will be combined with it.

The following example will help make things clear:

An account has a list of discount plans, among which, for example, 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 discount plan assigned directly to an account takes the highest precedence and defines how a discount plan added via an add-on product is combined with it. Then the discount plan that's assigned to the account via the add-on product, in turn, defines the way the discount plan that was added via the main product is combined with it.

If any discount plan applicable to a session is used up, the remaining discount plans are combined following the same rule.

The section below will demonstrate the ways to combine discount plans.

Ways to Combine Discount Plans

There are several ways that discount plans can be combined with each other. These are defined within the Discount Wizard and can be one of the following: **Never**, **Always**, **When discount lower than 100 is active** and **After reaching the last threshold limit**.

Never

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

Always

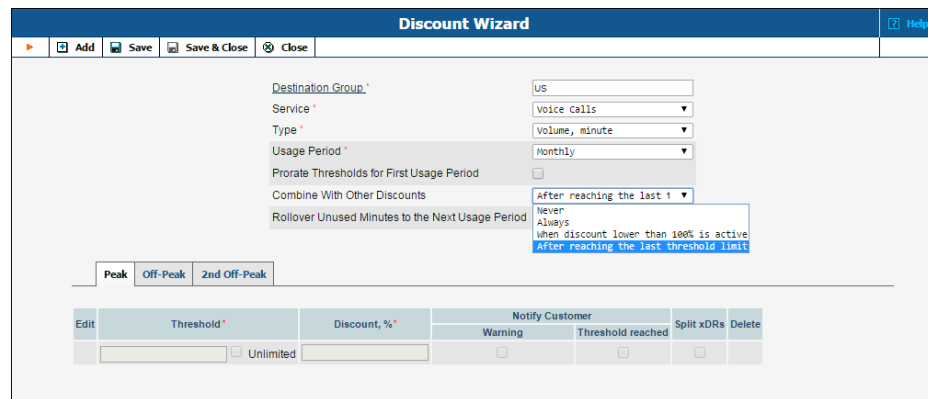
Using this way of combining, discounts are summed up to 100% (there is no money back for a consumed service). For example, two 30% discounts will result in a 60% final discount; 70% + 40% discounts will result in a 100% final discount.

When discount lower than 100 is active

If a higher precedence 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 precedence 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 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) 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 (SIP Product) that sets the rates – so that calls to the US&Canada cost 0.2 USD/min. The administrator assigns 2 discount plans to this account in the following order:

1. A 50% discount on calls to the US for the first 60 minutes.
2. A 100% discount on calls to US&Canada for the first 20 minutes.

Account Info / Retail Customer 'EasyCall Ltd'									
Account ID: 12126505550 Balance Control Subordinate									
Discount Plan: No volume discount									
Discount Plans									
History	Destination Group ↑	Volume Discount Plan	Precedence	Combined Discounts		Previous Threshold	Used/Remaining	Current Threshold	Expiration
				Peak Level	Dis...				
Voice Calls									
	US	USA Cheap	Add-On Pro...	50%	0 minute	0 minute	60 minute	2 days	
	US & Canada	US&Canad...	Add-On Pro...	100%	0 minute	0 minute	20 minute	2 days	

The account number 12126505550 makes the next calls:

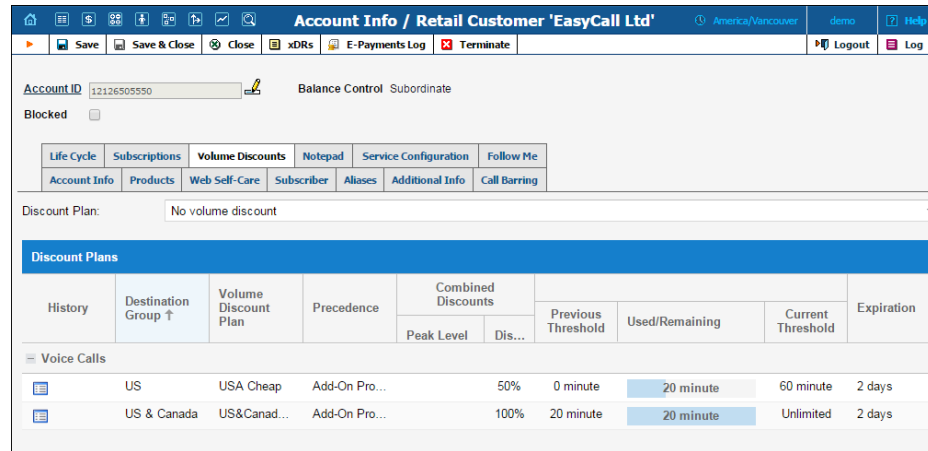
1. First, to Canada – for 15 minutes – it is charged using the “US&Canada” discount plan and comes to 0.0 USD.

Account Info / Retail Customer 'EasyCall Ltd'									
Account ID: 12126505550 Balance Control Subordinate									
Discount Plan: No volume discount									
Discount Plans									
History	Destination Group ↑	Volume Discount Plan	Precedence	Combined Discounts		Previous Threshold	Used/Remaining	Current Threshold	Expiration
				Peak Level	Dis...				
Voice Calls									
	US	USA Cheap	Add-On Pro...	50%	0 minute	0 minute	60 minute	2 days	
	US & Canada	US&Canad...	Add-On Pro...	100%	0 minute	15 minute	20 minute	2 days	

2. Then, to the US – for 20 minutes – it is charged using the “US” discount plan and comes to 2 USD (20 * 0.2 * 50%).

Account Info / Retail Customer 'EasyCall Ltd'									
Account ID: 12126505550 Balance Control Subordinate									
Discount Plan: No volume discount									
Discount Plans									
History	Destination Group ↑	Volume Discount Plan	Precedence	Combined Discounts		Previous Threshold	Used/Remaining	Current Threshold	Expiration
				Peak Level	Dis...				
Voice Calls									
	US	USA Cheap	Add-On Pro...	50%	0 minute	20 minute	60 minute	2 days	
	US & Canada	US&Canad...	Add-On Pro...	100%	0 minute	15 minute	20 minute	2 days	

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



The screenshot shows a web application interface for 'Account Info / Retail Customer 'EasyCall Ltd''. It features a navigation bar with options like 'Save', 'Save & Close', 'Close', 'xDRs', 'E-Payments Log', and 'Terminate'. Below the navigation bar, there are tabs for 'Life Cycle', 'Subscriptions', 'Volume Discounts', 'Notepad', 'Service Configuration', and 'Follow Me'. The 'Volume Discounts' tab is active, displaying a table of discount plans. The table has columns for History, Destination Group, Volume Discount Plan, Precedence, Combined Discounts (Peak Level, Dis...), Previous Threshold, Used/Remaining, Current Threshold, and Expiration. Two plans are listed under 'Voice Calls': 'US' with a 50% discount and 'US & Canada' with a 100% discount.

History	Destination Group ↑	Volume Discount Plan	Precedence	Combined Discounts		Previous Threshold	Used/Remaining	Current Threshold	Expiration
				Peak Level	Dis...				
[-]	US	USA Cheap	Add-On Pro...	50%	0 minute		20 minute	60 minute	2 days
[-]	US & Canada	US&Canad...	Add-On Pro...	100%	20 minute		20 minute	Unlimited	2 days

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.

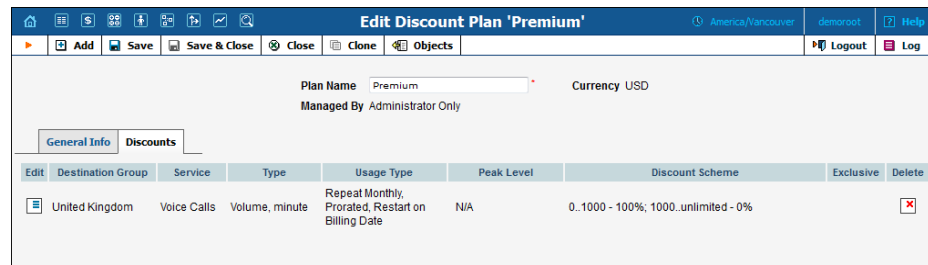
Prorated thresholds in volume discount plans

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:

- **Free minutes** – A volume discount plan with a monthly usage period is created with a 100% discount for up to 1000 minutes, and after 1000 minutes a regular rate is applied.
 - If the volume discount plan is assigned on October 1st, then the threshold is 1000 minutes.
 - If the volume discount plan is assigned on October 20th, then the threshold becomes 387 minutes (since 12 days remain in October).
 - If the volume discount plan is assigned on October 30th, then the threshold is 65 minutes (since 2 days remain in October).

The threshold is then 1000 minutes for the following month and thereafter.



Edit	Destination Group	Service	Type	Usage Type	Peak Level	Discount Scheme	Exclusive	Delete
	United Kingdom	Voice Calls	Volume, minute	Repeat Monthly, Prorated, Restart on Billing Date	N/A	0..1000 - 100%; 1000..unlimited - 0%		

- 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 the volume discount plan is assigned 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.



Edit	Destination Group	Service	Type	Usage Type	Peak Level	Discount Scheme	Exclusive	Delete
	Eastern Europe	Voice Calls	Volume, minute	Repeat Weekly, Prorated	N/A	0..100 - 0%; 100..200 - 10%; 200..unlimited - 20%		

- Change of Volume Discount Plan** – The “**Premium**” volume discount plan is created with a 100% discount for the first 1000 minutes and after 1000 minutes a regular rate is then applied. A “**Standard**” volume discount plan is created with a 100% discount for to the first 500 minutes and after that a regular rate is applied. The volume discount plan has a monthly usage period. If the “**Premium**” volume discount plan is assigned on October 8th, then the threshold is 774 minutes (since 24 days remain). If the volume discount plan is changed to “**Standard**” on October 21st, the threshold is 177 minutes (since 11 days remain).

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

When an end user switches from one discount plan to a new one with a similar configuration that is applicable to the same destination group and has been assigned to an account at the same level (i.e. via the add-on product of the same precedence level or included in the account’s main product, etc.), the volume of consumed discounted service is transferred to the new discount plan.

This means that the counters of the previous discount plan are preserved for use with the new one. For example, a user has 50 free domestic SMS messages after sending 25 SMS with no discount per month. He has sent 25 SMS already but wants to upgrade to 100 domestic SMS messages this month. He signs up for a new SMS bundle and since he has already sent 25 SMS messages, 100 free domestic SMS will be available for him immediately after the discount plan change.

As mentioned above, to facilitate the transfer of consumed discounted service volume, the new discount configuration must have the following parameters identical to the previous configuration:

- A destination group that belongs to the same destination group set (e.g. both discount plans apply to the US & Canada destination group from the North America destination group set);
- Service type (e.g. for both discounts the service type is defined as Voice Calls);
- Threshold type (e.g. the provided discount is based on the session cost);
- 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).

When assigning a new discount plan to a customer / account the following conditions must be met:

- Both discount plans must have *the same priority level* (e.g. the previous and new discount plans are associated with the main product or the add-on products at the *same* precedence level);
- The discount plan change must be done as a single operation;

NOTE: When changing discount plans for a customer or account via the web interface, remove the old discount plan, assign a new one and click **Save**.

When discount plans for a customer or account are changed using the API, the following methods are used:

- To change discount plans included in the main product use the `i_vd_plan` property in the `Update_product` method;
- To change discount plans for a particular customer use the `i_vd_plan` property in the `Update_customer` method;
- To change 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 [PortaSwitch External Interfaces Guide](#) for more information about API methods.

When changing a discount plan for a customer / account, you may also want to define a different threshold structure in a new discount. This will allow you to extend / reduce existing discount volumes for your customers.

For example, a user signs up for 10 GB of Internet traffic quota per month. He uses up 9 GB in 10 days and wants to sign up for 20 GB traffic quota during this month. He checks that such a product is available and switches over to it. Now the user has the 20 GB product, but since 9 GB are already used, only 11 GB are available for the user this month.

If you want to reduce the existing discount volume for your customers, assign a new discount plan that provides less volume of discounted service than the previous one. Then if the volume of already consumed discount is higher than the total volume of new discount, the newly assigned discount plan will show the status **Used in Full** until it is reset. For example, the user from the previous example 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 7 GB Internet traffic quota available.

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

If any of the above-described conditions are not met, the volume of consumed discounted service will not be transferred. So for instance, the user who had 50 free domestic SMS messages and now wants 100 global SMS messages per month switches over to a new product and receives 100 global SMS per month, regardless of whether he used any domestic SMS or not.

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.

Off-peak periods and volume discounts

By default, volume discounts 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 discount plan will be modified, based on the time when the service is used.

Customer notification before reaching a volume discount threshold

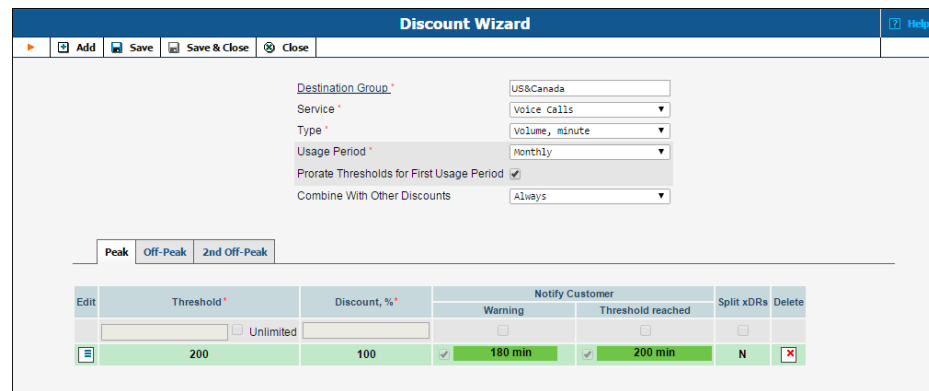
You can notify your customers that they are reaching the threshold for when a current discount stops being provided to them.

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 'Discount Wizard' interface. It includes a toolbar with 'Add', 'Save', 'Save & Close', and 'Close' buttons. The main configuration area has the following fields:

- Destination Group: US&Canada
- Service: Voice calls
- Type: Volume, minute
- Usage Period: Monthly
- Prorate Thresholds for First Usage Period:
- Combine With Other Discounts: Always

Below these fields are tabs for 'Peak', 'Off-Peak', and '2nd Off-Peak'. A table below the tabs shows the configuration for a single discount plan:

Edit	Threshold *	Discount, % *	Notify Customer		Split xDRs	Delete
			Warning	Threshold reached		
<input type="checkbox"/>	<input type="checkbox"/> Unlimited	<input type="text"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>
<input checked="" type="checkbox"/>	200	100	<input checked="" type="checkbox"/> 180 min	<input checked="" type="checkbox"/> 200 min	N	<input checked="" type="checkbox"/>

Then he assigns this 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 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 xDR

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 normal 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.

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 services. The bundle promotions feature allows you to offer promotional discounts based on monthly amounts spent and apply charges based on usage minimum.

The key features of the bundle promotions include the following:

- Promotional discounts based on the amount of money spent for a service within the billing period. For example, if you wish to offer 10 percent 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.

Edit Bundle Promotion '10% voice calls disc. on \$1000'

Name: 10% voice calls disc. on \$1000
 Description: 10% discount for voice calls on minimum of \$1000 spent
 Managed By: Administrator Only
 Currency: USD - US Dollar
 Destination Group Set: Corporate - World

Analyze		Criteria			Apply To		Delete
Service	Destination Group	Apply	Structure	Invoice Comment	Service	Subscription	
Conferencing		Credit			Whole Bill		
Voice Calls	World	Credit	0..1000 USD - 0 % 1000..Unlimited USD - 10 %		Whole Bill		X

- Promotional discounts are given for the actual service usage of the promoted service. For example, if you wish to provide \$10 off for SMS 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 SMS, will receive an SMS discount of \$8 of a possible discount of \$10.

Edit Bundle Promotion 'Credit on SMS'

Name: Credit on SMS
 Description: Spend more than \$50 on voice calls and receive up to \$10 discount on SMS
 Managed By: Administrator Only
 Currency: USD - US Dollar
 Destination Group Set: Retail

Analyze		Criteria			Apply To		Delete
Service	Destination Group	Apply	Structure	Invoice Comment	Service	Subscription	
Conferencing		Credit			Whole Bill		
Voice Calls		Credit	0..50 USD - 0 USD 50..Unlimited USD - 10 USD		Messaging Service		X

- 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.

Edit Bundle Promotion 'Corporate Plus'

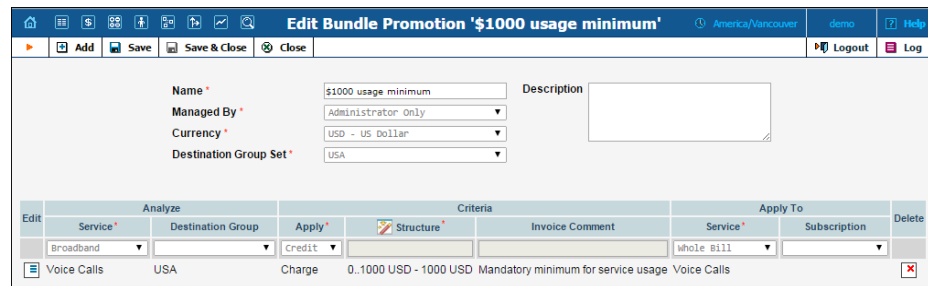
Name: Corporate Plus
 Description: up to \$50 - no discount, up to \$100 - 10%, over - 20% discount
 Managed By: Administrator Only
 Currency: USD - US Dollar
 Destination Group Set: Corporate - World

Analyze		Criteria			Apply To		Delete
Service	Destination Group	Apply	Structure	Invoice Comment	Service	Subscription	
Conferencing		Credit			Whole Bill		
Voice Calls	Asia	Credit	0..50 USD - 0 % 50..100 USD - 10 % 100..Unlimited USD - 20 %		Voice Calls		X

- Mandatory minimum for service usage. Service providers 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.

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, and as a result, the deficient \$200 amount is added to his invoice so that it equals \$1000.



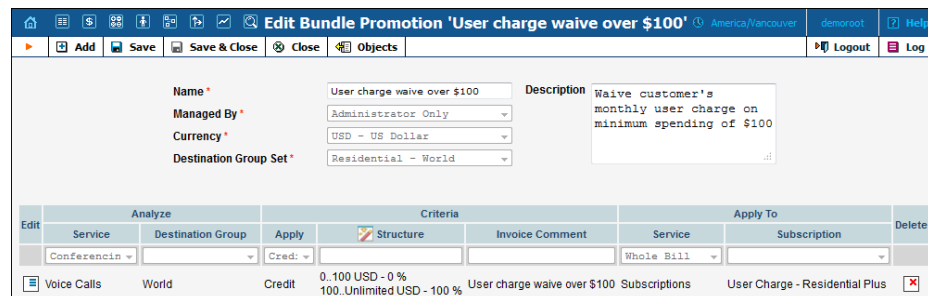
The screenshot shows the 'Edit Bundle Promotion' interface for a '\$1000 usage minimum' promotion. The form includes fields for Name, Managed By, Currency, and Destination Group Set. Below the form is a table with columns for Analyze, Criteria, and Apply To. The table contains one entry for 'Voice Calls' with a 'Mandatory minimum for service usage' criteria.

Analyze		Criteria		Apply To		Delete	
Service	Destination Group	Apply	Structure	Invoice Comment	Service	Subscription	
Voice Calls	USA	Charge	0..1000 USD - 1000 USD	Mandatory minimum for service usage	Voice Calls		X

This approach means the revenue stream is more predictable for service providers and allows them to plan their expenses more efficiently and become more financially stable.

- 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 the 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.

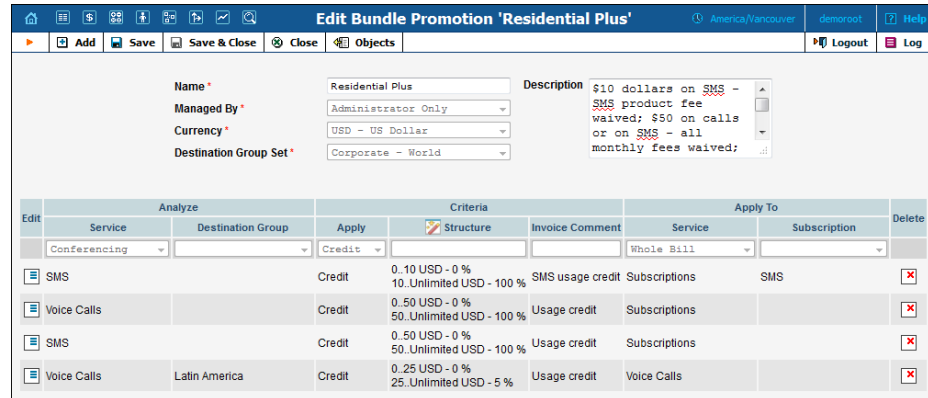


The screenshot shows the 'Edit Bundle Promotion' interface for a 'User charge waive over \$100' promotion. The form includes fields for Name, Managed By, Currency, and Destination Group Set. Below the form is a table with columns for Analyze, Criteria, and Apply To. The table contains one entry for 'User Charge - Residential Plus' with a 'User charge waive over \$100' criteria.

Analyze		Criteria		Apply To		Delete	
Service	Destination Group	Apply	Structure	Invoice Comment	Service	Subscription	
User Charge - Residential Plus	World	Credit	0..100 USD - 0 % 100.Unlimited USD - 100 %	User charge waive over \$100	Subscriptions		X

- Multiple promotions in the same promotion bundle. For example, your promotion bundle could state: “Spend \$10 dollars on SMS 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 South America and you will get a 5% discount on all your voice calls.”



- The discount can either be a percentage of the charged amount in the 0 – 100 range (a 0% discount means “no discount applied,” while a 100% discount means “all charges waived”), or a fixed value.

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 a credit or a charge.

The promotion will be applied at the end of the customer’s billing period – when all customer service charges and balance adjustments have been calculated – and will be reflected in a separate line(s) on the invoice.

Let’s look at the following example:

You apply the promotion: “10 percent off if the customer spends over \$1000 on voice calls” (see the first example above).
At the end of the month a customer has spent \$1200 on phone calls.

When the customer’s billing period closes, an extra xDR will be created and the customer will see a line with a \$120 **credit** on his invoice. The description for this transaction will be **10% (\$1,200)** so it will be easy to understand what it relates to. The invoice total will be \$1,080.

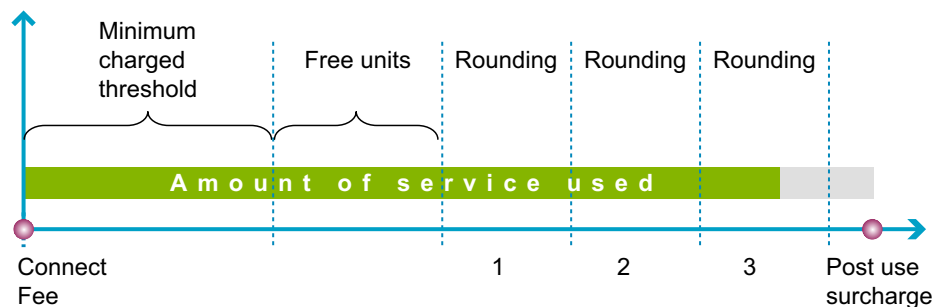
The differences between bundle promotions and volume discounts

Unlike our volume-based discount, 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 the 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

Quantity Billing Parameters

Quantity-based services (such as messaging) 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:

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

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:

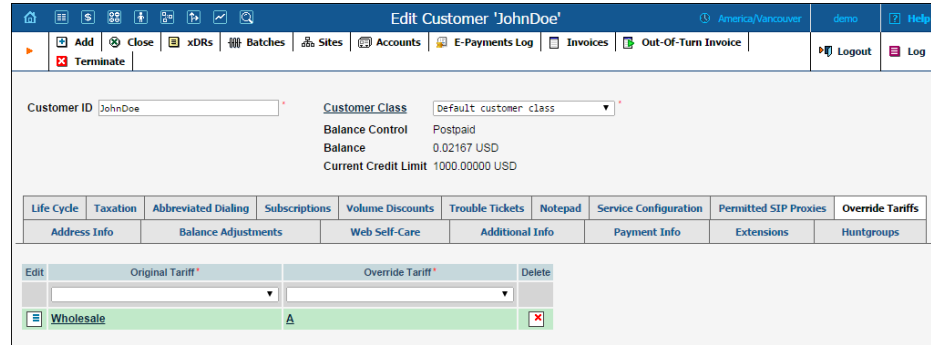
- The measurement unit is a byte, the billing unit is a kilobyte, and the billing ratio (how many measurement units make up one billing unit) is 1,024.
- The rate is defined as follows: minimum threshold – 10,240, rounding – 1,024, and the initial and next prices are the same – \$0.02. (Prices are defined 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 8 rounding increments above the minimum threshold, and so the total charged amount is $10,240 * 0.02 / 1,024 + 8 * 1,024 * 0.02 / 1,024 = \0.36 .


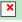
Override Tariff

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.

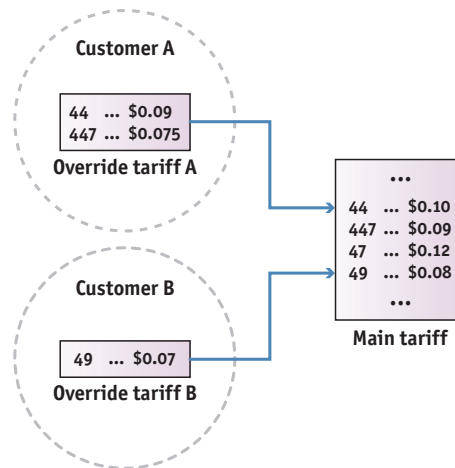


Life Cycle	Taxation	Abbreviated Dialing	Subscriptions	Volume Discounts	Trouble Tickets	Notepad	Service Configuration	Permitted SIP Proxies	Override Tariffs
Address Info		Balance Adjustments		Web Self-Care		Additional Info		Payment Info	
Extensions		Huntgroups							

Edit	Original Tariff *	Override Tariff *	Delete
	Wholesale	A	

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).

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 react 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 causes inconvenience to 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 as to whether the currently available funds are enough to cover a service (or how much of the 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 money 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 only be working with 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, in order 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, which is used when it is not possible to tell how much a customer will spend in a particular session, so that additional funds have to 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 the available funds to cover the next time interval (e.g. 5 minutes). As the end of this interval comes near, the system attempts to reserve funds to cover the next interval, and so on. 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), this will immediately affect the call – either it will be disconnected sooner, or the maximum allowed call time will increase.

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).

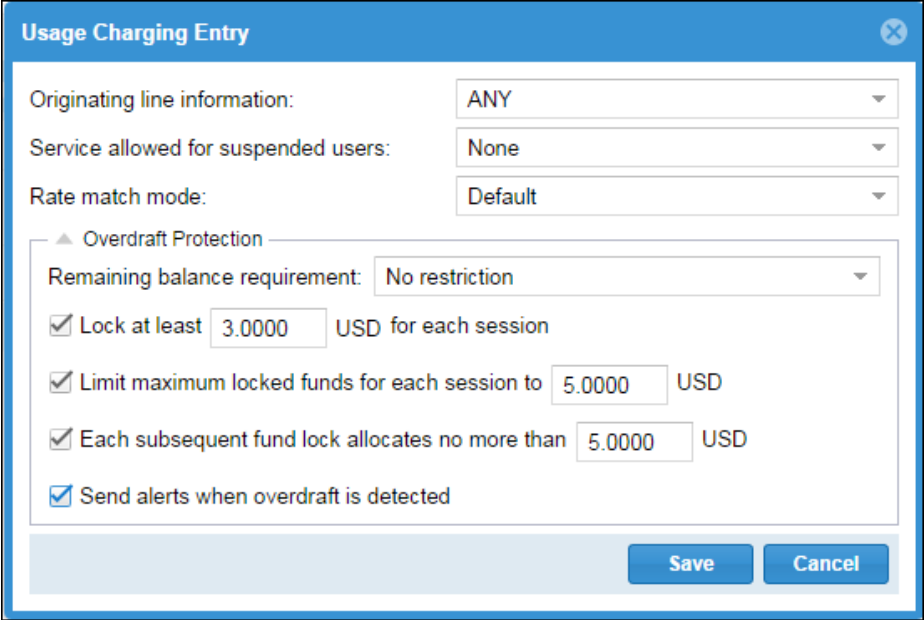
- Since the user now has available funds, another 5 minutes of call time are authorized. The total amount of locked funds is now \$8.40 (11).
- Finally, the user disconnects, and is charged \$7.80 for the full call duration of 26 minutes. Thus his balance is now \$3.20 (12).

Limiting the maximum fund amount locking

Two situations 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 US or Europe (in which case just 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 up, and will not be able to 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 which does not support dynamic fund locking, then all the 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 image shows a dialog box titled "Usage Charging Entry" with a close button (X) in the top right corner. The dialog contains several configuration options:

- Originating line information: ANY (dropdown)
- Service allowed for suspended users: None (dropdown)
- Rate match mode: Default (dropdown)
- Overdraft Protection (expanded section):
 - Remaining balance requirement: No restriction (dropdown)
 - Lock at least 3.0000 USD for each session
 - Limit maximum locked funds for each session to 5.0000 USD
 - Each subsequent fund lock allocates no more than 5.0000 USD
 - Send alerts when overdraft is detected

At the bottom right, there are "Save" and "Cancel" buttons.

In this case each request will lock up 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 authorization every three minutes.

Legacy equipment and backward compatibility

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. Currently, real-time overdraft protection is fully supported by the PortaSIP® call processing engine.

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 previously available “fraud protection” feature in PortaBilling® (where a debit account is only allowed to make one simultaneous call)?

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

The screenshot shows the 'New Product' configuration page in PortaBilling®. The 'Product Name' is 'EasyCall'. The 'Overdraft Protection' checkbox is checked. The 'Breakage' is set to 0.00 USD. The 'Service Used At' table shows 'PortaSIP' for 'Voice Calls'.

Edit	Log	Service Used At			Charges Applied		Delete
		Node	Access Code	Service	Tariff	Advanced Config	
		PortaSIP	Any code	Voice Calls	Prepaid Cards		

Overdraft Protection:

Breakage: 0.00 USD

- When overdraft protection is **disabled**, it means that there are effectively no “locked” funds. When a locking of funds is requested, they are done separately for each session, and do not affect other sessions. So there can be several concurrent sessions, and each can potentially use up all of 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 amount, 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.

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 RADIUS 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 RADIUS 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 specifies the “original” (incorrect) tariff and the correct tariff, and then narrows the set of xDRs to be processed by time interval and specific customer.

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).

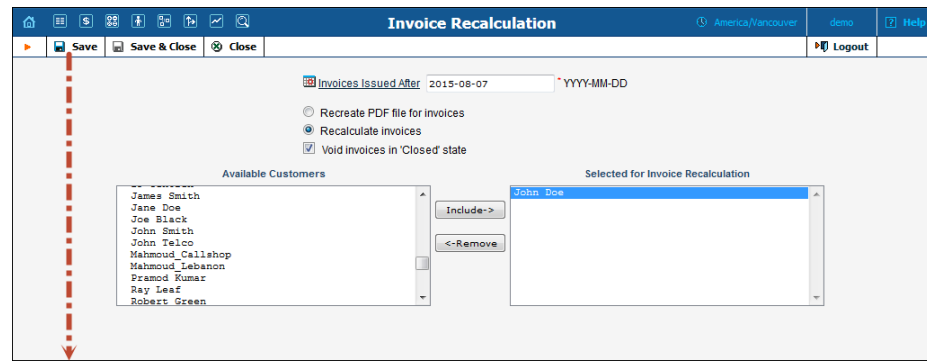
Correction mode

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 recalculation process on the **Invoices** page.

The screenshot shows the 'Invoices' page in the PortaBilling web interface. The page title is 'Invoices' and it shows a search form for invoices. The search criteria include Invoice No., Customer (set to ALL), From date (2015-07-07), To date (2015-08-07), and a checkbox for 'Include Void Invoices'. A red dashed arrow points to the 'Invoice Recalculation' button in the top navigation bar.



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, this process should always be run from a specific moment in the past (when the error occurred) to the present moment. In this case, all discount counters will be rolled back before recalculation actually starts, and then updated with each re-rated call.

Subscriptions

The **Subscriptions** module allows you to charge your customers periodic fees 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 here. By defining a

subscription plan you can specify various parameters within this plan that will have an effect on the customer balance. Every subscription plan includes the following:

- Name and description – Used by your administrators to better manage various subscription plans.
- Subscription plan name visible to end user – A clear subscription name shown to end users on invoices and on their self-care interfaces.
- Activation mode – Specifies when the subscription is active and charges start to be assessed. Typically, a subscription is activated immediately upon its creation (or on the start date, if that has been specified). PortaBilling® supports an additional mode of subscription activation based on the account's first use. This allows you to avoid problems when there is a delay between the time a customer signs up for a service and the time he is actually able to use it (e.g. he uses online signup to purchase a residential SIP service, but his IP phone will not be delivered by FedEx until five days later).
- Activation fee – The amount to be charged when the subscription becomes active.
- 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 he keeps paying monthly fees until he has paid back the cost of the phone). If the subscription service is canceled earlier than the interval specified, the customer will be charged a cancellation penalty.
- Rounding – To avoid unusual subscription charges like \$1.26789, you can specify a rounding pattern for them (similarly to what you did with ratings in tariffs). For example, the rounding pattern XXXXX.XX000 will round up all subscription charges to whole cents.
- Of course, every subscription also includes a definition of periodic 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:

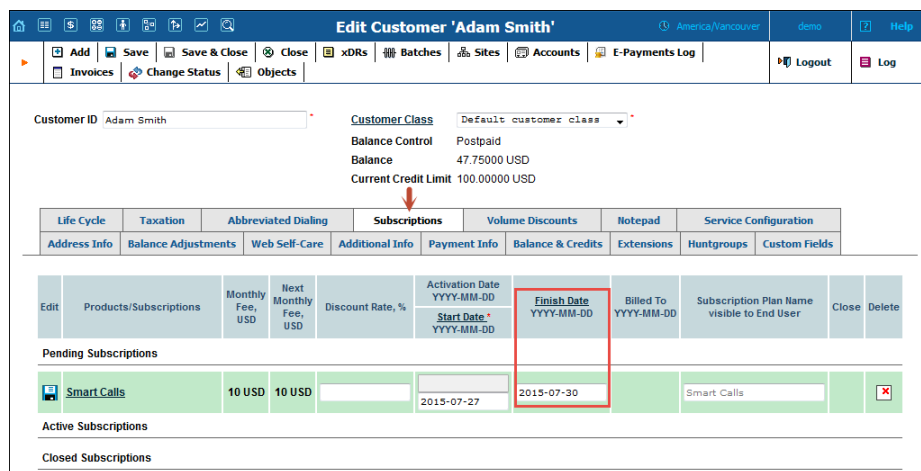
- Activation fee
- Cancellation fee
- Periodic fee

Activation fee

An activation fee is automatically activated when a subscription start time begins, with one exception. Namely, in the subscription plan’s properties you may specify that the subscription plan will become active only when the account is used for the first time. In this case, the subscription is activated either on its start date or the account’s date of first use, whichever comes later.

Cancellation fee

A subscription which is not yet active can simply be deleted (for instance, a customer signed up for a new service beginning next month, but then changed his mind). A subscription which is already active cannot be deleted. When the subscription is canceled, the cancellation date is entered as the end date of the subscription. Typically, the cancellation date will be the day on which cancellation took 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 subscribed to the service on the 3rd and cancelled it on the 7th, he will still be charged for five days of service.

Periodic fee

A periodic fee defines the recurring charges applied to a customer for service usage. Periodic subscription fees are pro-rated depending on the actual duration of the service used.

Consider the following examples related to the same subscription plan with a monthly periodic fee of \$9.99:

- Customer A, with a monthly billing period, activates the subscription on April 12th. On May 1st he is charged \$6.33 because

he uses the service for 19 days (12-30) and $(\$9.99/30) * 19 = \6.33 .

- Customer B, with a monthly billing period, activates the subscription on April 12th and cancels it on the 25th. On May 1st he is charged \$4.66 because he uses the service for 14 days (12- 25) and $(\$9.99/30) * 14 = \4.66 .

These types of subscription fees are reflected on the invoice as extra invoice lines.

Periodic Fees Configuration

Since your customers may use different billing periods, PortaBilling® gives you 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 increased maintenance required by customers who have a shorter billing period justifies increasing rates.

For instance, PortaBilling® allows you to define your monthly rate as \$19.99, your bi-weekly rate as \$10.99, weekly rate as \$6.99 and daily rate as \$1.99 for the 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 \$19.99 for the next nine months, and \$12.99 thereafter; or something even more complex.

Incompletely used promotional periods

The promotional period in a subscription includes each billing period during which a service was used, regardless of the period of days used in a particular billing period.

For example, you offer a promotional period for three months. Customer John Doe subscribes for the promotion on July 15th. John Doe will be charged for the second half of July (16 days) and for the two months following it (August and September) according to the promotional rate. Starting from the fourth month (October), the default rate will be 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 the two 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 some 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.

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.

Changing the subscription plan for a product

You may change the subscription plan for a given product at any time; note, however, that this will affect all existing accounts that have this product. Once you have selected a new subscription plan for a product, this change takes effect immediately (i.e. 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) will be changed to the subscription plan assigned to product B.
- Optional subscriptions (assigned to the account directly) will 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 subscription fees ensures that, at the end of the billing period, the customer is charged until the end of the Nth full billing period following the current one. So if we charge one month in advance and the current monthly billing period ends on April 10th, the customer will be charged until May 10th, and his invoice produced on April 11th will contain charges for calls made prior to April 10th plus subscription charges covering the period from April 11th to May 10th.

If a subscription that is charged in advance is not active during the whole billing period, it is pro-rated in the same way that post-charged subscriptions are. For instance, if a customer with a monthly billing period activates a \$30 subscription plan charged one month in advance on April 15th, he will only be charged \$15 ($\30 of monthly fee / 30 days in the month * 15 days of the actual service usage). The customer's invoice for April will include a \$15 subscription charge for the period from April 15th to April 30th and a \$30 subscription charge for the period from May 1st to May 31st (in addition to charges for any calls made in April).

If a customer terminates a subscription prior to the end date of a period that has already been charged (e.g. the customer in the example above closes his account on April 20th), then a refund is issued for the "unused" subscription time when the billing period is closed.

When periodic fees are modified during a current billing period, a subscription with fees charged in advance will cost the same until the end of that current billing period. In other words, there will be no recalculation. The charges for the next billing period 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. But if the subscription fee changes on April 25th and becomes \$8, the customer's invoice (produced on April 11th) contains charges for all calls made prior to April 10th at the old fee 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) will be applied to the new billing period, which will cover the period from May 11th to

June 10th, and the customer will continue to be charged \$8 for subsequent billing periods.

However, if a customer terminates a subscription prior to the end date of a period that has already been charged (e.g. the customer in the example above closes his account on May 1st and the subscription fee changes from \$10 to \$8 on April 25th), then a refund will be issued for the “unused” subscription time when the billing period is closed, but the new subscription fee will be used for recalculating (i.e. \$8 not \$10).

This can improve a service provider’s cash-flow and minimize potential loss if a customer leaves without paying his final bill.

Subscriptions charged in advance to prepaid customers

In most cases subscriptions that are charged in advance are assigned to prepaid customers. Therefore, an administrator must be sure that customers have sufficient funds to cover their subscriptions and when they don’t, the customers’ service must be suspended.

For this there is an option of automatically suspending prepaid customers’ service when they do not have sufficient funds 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. The amount of available funds is now \$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 was 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 one for the refund amount (-\$4).

Thus, prepaid customers can only consume services when they have sufficient available funds to cover their next subscription charges.

Auto-charging for subscriptions and service usage balance

Another popular business model of applying subscription charges is to automatically deduct them from customers' credit cards. Thereby the service provider maintains the "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 available funds are consumed by subscription charges.

To automatically charge customers for subscriptions, configure a customer class as follows:

- Enable the **Separate Invoice for Recurring Fees** option;
- Enable the **Auto-charge invoice balance to credit card** option;
- Specify the invoice template;
- Specify the grace period, for example, 5 days;
- Specify when to charge the credit card for the unpaid invoice, 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.

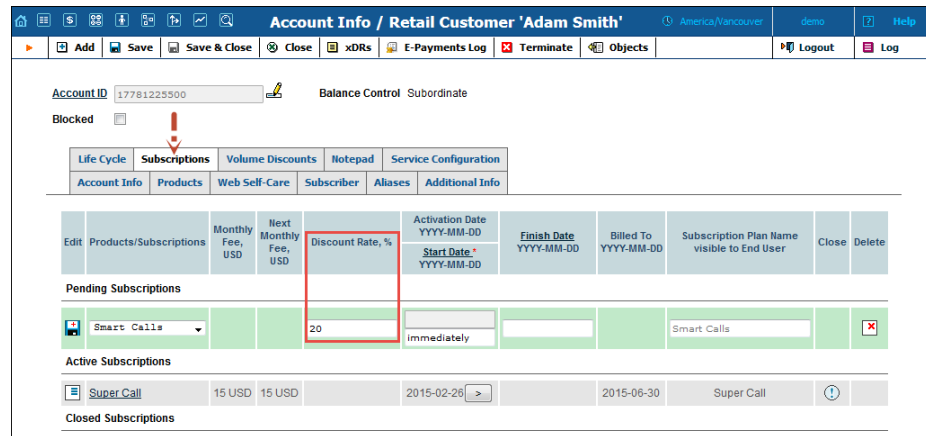
With auto-charging for subscriptions, there is improved cash flow since there is pre-defined recurring revenue.

Customers are better able to control their available funds since these funds are only used for services.

Subscription 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. So 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 will be 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.

Prepay 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 prepay 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).

Edit Subscription Plan 'IPTV Basic' Europe/Prague demo Help

Subscription Plan Name: IPTV Basic Currency: USD
 Subscription Plan Name visible to End User: IPTV Basic Managed By: Administrator Only

Percentage
 Fixed Amount

Edit	months	Standard Fee, USD / months	Discount, %	Discount Fee, USD / months	Delete
<input type="checkbox"/>	6	119.94 USD / 6 months	10	107.946 USD / 6 months	<input type="checkbox"/>
<input type="checkbox"/>	12	239.88 USD / 12 months	20	191.904 USD / 12 months	<input type="checkbox"/>

Edit Customer 'John Doe' Europe/Prague demo Help

Customer ID: John Doe Customer Class: Default customer class
 Balance Control: Postpaid Balance: 109.94291 USD Current Credit Limit: 500.00000 USD

Edit	Subscription	Current Fee	Next Fee	Discount Rate, %	Activation Date YYYY-MM-DD Start Date * YYYY-MM-DD	Finish Date YYYY-MM-DD	Billed To YYYY-MM-DD	Subscription Plan Name visible to End User	Close	Delete
Pending Subscriptions										
<input type="checkbox"/>	IPTV Basic	191.904 U			Immediately			IPTV Basic	<input type="checkbox"/>	<input type="checkbox"/>
Active Subscriptions										
Closed Subscriptions										

Then when assigning a subscription to a customer or an account, a corresponding prepay 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
 $(\$19.99 * 6 - 10\% = \$119.94 - \$11.99 = \$107.95);$
- 20% discount for 12 months
 $(\$19.99 * 12 - 20\% = \$239.88 - \$47.98 = \$191.90).$

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.)

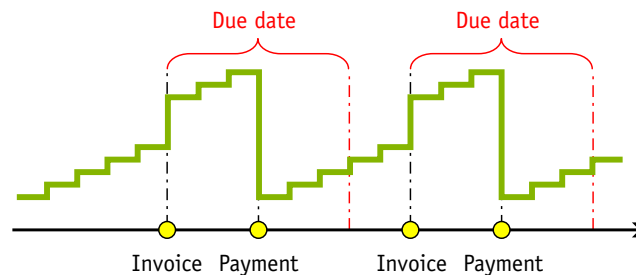
Thus, customers have the option to choose how to pay for their services. When paying months in advance, they receive a discount on services and, as a result, save money.

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.

Invoicing


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.




- Bi-weekly – 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, his 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 use a special “Do not create invoice” option for a customer. In this case, PortaBilling® will only produce CDR summaries, but no actual invoices will be generated.

Please consult the [PortaBilling Templates Guide](#) for more information about creating and managing invoice templates.

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 Info** tab).

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. 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.

You can see the customer’s current “unallocated payment” status on the **Payment Info** 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.

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.

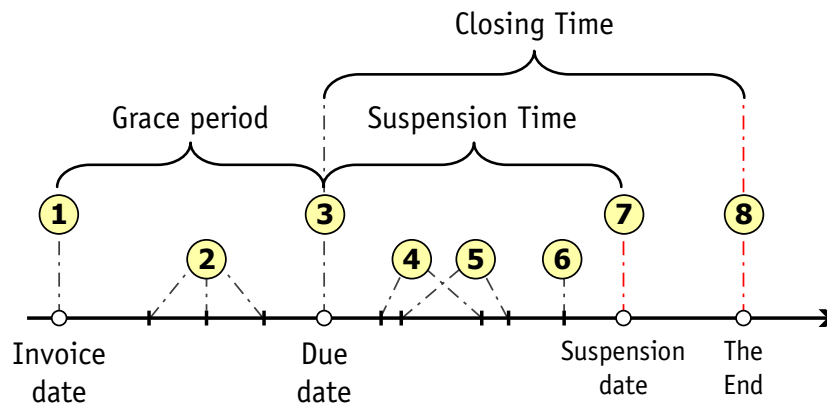
Customer account during collection

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 correct 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 for regeneration (if, for example, during an invoice review the administrator finds something wrong with the invoice template requiring him to change it). Only after that step are the invoices delivered via email to your customers.



This functionality can be used for the following situations:

- **Billing of premium customers** – If you have premium customers you naturally want to provide them with high quality services. The invoices are no exception. You want to make sure that all charges are included and that the invoice total is calculated correctly.

- **Billing of a new service** – An administrator creates a new service with new ratings, tariffs, products and invoice templates. If he is not sure whether all the new charges are correctly calculated he can activate the “send invoice after review and approval by administrator” option for a specific customer class. At the end of the billing period the invoices are created and marked as “needs review,” but are not yet sent to customers and not available for download from the web interface. Only after the administrator checks and approves the invoices they are delivered to customers via email.
- **Changes in templates** – If, during the invoice review an administrator discovers a problem with a template and a .pdf file also does not appear to be correct, he changes the template. Then, on the **Invoices Review** page he selects all of the “under review” invoices that were created using this template and regenerates correct ones. New .pdf's are then generated and 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 2013 with an invoice total of \$150, notices an 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.

Midterm Invoices

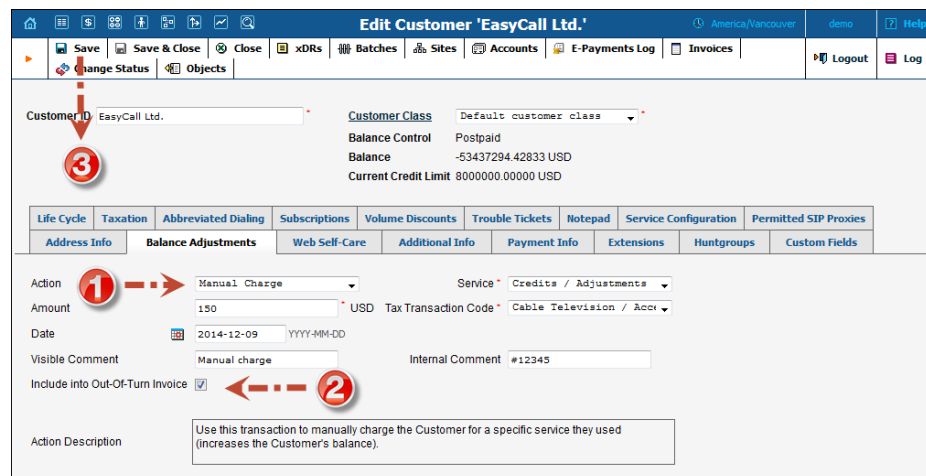
It might occur that in the middle of a billing period, a customer wants to pay for services consumed and does not want to wait until the end to receive an invoice. For such customers, administrators may now generate a midterm invoice.

Let's assume that a customer with a monthly billing period wishes to pay for services consumed up to the 13th of May and informs the administrator about this. On the 14th of May the administrator generates a midterm invoice for this customer, which covers the period from the 1st of May until the 13th of May. The invoice is then immediately sent to the customer. The regular invoice that's generated at the end of the billing period will cover the rest of the period – from the 14th of May until the 31st of May.

Thus, midterm invoice functionality will help 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 Invoice Generation

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 'Edit Customer' interface for 'EasyCall Ltd.'. The form is titled 'Manual Charge' and includes the following fields and options:

- Action:** Manual Charge (highlighted with a red arrow and number 1)
- Amount:** 150 USD
- Date:** 2014-12-09 (YYYY-MM-DD)
- Visible Comment:** Manual charge
- Internal Comment:** #12345
- Include into Out-Of-Turn Invoice:** (highlighted with a red arrow and number 2)
- Action Description:** Use this transaction to manually charge the Customer for a specific service they used (increases the Customer's balance).

Customer information displayed includes: Customer Class: Default customer class, Balance Control: Postpaid, Balance: -53437294.42833 USD, and Current Credit Limit: 8000000.00000 USD. A red arrow and number 3 point to the 'Customer' dropdown menu.

For example, if your customer visits your office to buy a new IP phone, your clerk will locate the customer's account, issue an invoice covering the

cost of the phone, take the money for the phone and give the phone and the invoice to the customer.

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.

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 in the template, 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.

Rate	xDR	Invoice
\$0.12 (\$0.10 + 20%)	15 min	Services \$1.50
	\$1.80	Tax \$0.30
		Total \$1.80

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”.

Invoice templates

When creating a new invoice template, you will choose if this template is to be used to produce tax-inclusive or tax-exclusive invoices. If you choose the tax-inclusive option and specify the actual tax rates, the tax calculation formulas will automatically be entered in the post-processing rules in fields like Tax1, Tax2, and so on.

Tax Name	Rate	Delete
Tax 1	7	X
Tax 2	7.5	X

During invoice generation, PortaBilling® computes the tax amount based on the invoice total and the tax rate you have specified, and this information is displayed to the customer.

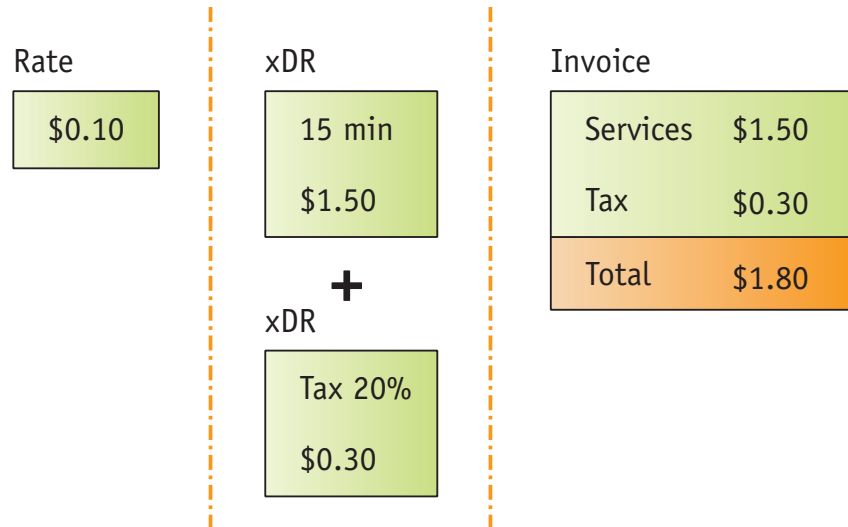

Back-calculation formulas support applying multiple types of taxes simultaneously. You simply need to enter the total tax rate and the rate of the specific tax. Our example above illustrates taxation for British Columbia (Canada), where there is 7% GST and 7.5% PST.

If you want to adjust the tax rate for an existing template or perform some other advanced modifications to the tax calculations, you can edit the contents of the post-processing rule formula directly. For more details about it and the **backcalcTax** function in particular, see the [PortaBilling® Templates Guide](#) manual.

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. In this case, the invoice template does not perform any tax-related calculations; it simply presents the sum total of all "tax" and other transactions.

NOTE: Applying an invoice template designed for the inclusive method to a customer who is charged by the exclusive method will most certainly produce incorrect calculations.



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. Each customer is assigned his own plug-in type, which defines a different module to be used for calculating taxes. Currently, PortaBilling® supports plug-ins for VAT (fixed percentage), custom taxation, the EZtax® (formerly known as BillSoft®) suite and the SureTax® taxation plug-in. The latter two enable correct calculation of US state, country and city taxes, as well as special items such as federal USF, etc.

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.

Custom taxation

This taxation plug-in is similar to VAT and is specifically designed for those who need to charge INT and UFS fees. 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 (formerly known as BillSoft) 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 (formerly known as BillSoft) 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:

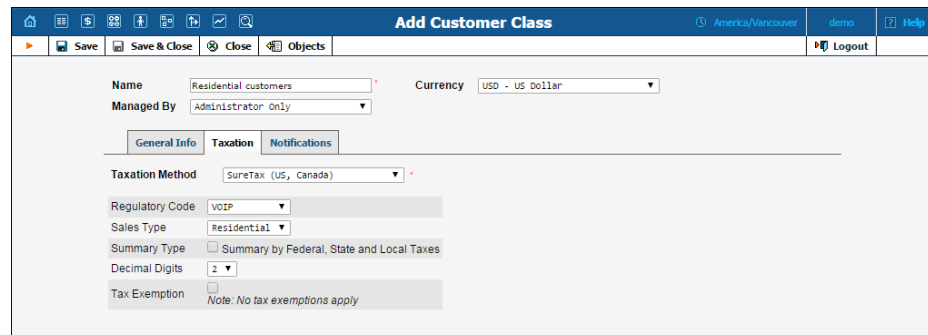
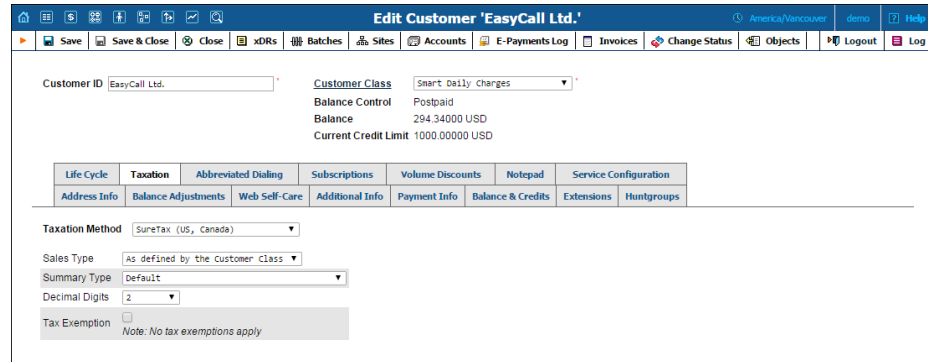
- Installation to the system's default folder.
- Installation to a separate folder for each of the system's environments.

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.

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 voice call services and subscriptions. SureTax (as a taxation method) can be defined within the customer class or for the individual customer.

When you select SureTax as a taxation plug-in, you can define such parameters as:

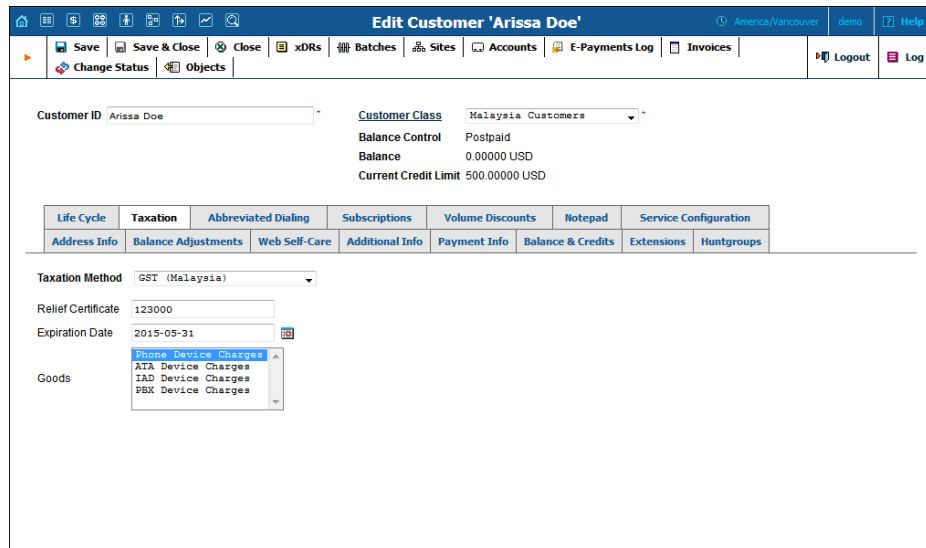
- type of customer;
- summary type;
- type of tax exemption.

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.



Edit Customer 'Arissa Doe'

Customer ID: Arissa Doe | Customer Class: Malaysia Customers

Balance Control: Postpaid
Balance: 0.00000 USD
Current Credit Limit: 500.00000 USD

Life Cycle	Taxation	Abbreviated Dialing	Subscriptions	Volume Discounts	Notepad	Service Configuration	
Address Info	Balance Adjustments	Web Self-Care	Additional Info	Payment Info	Balance & Credits	Extensions	Huntgroups

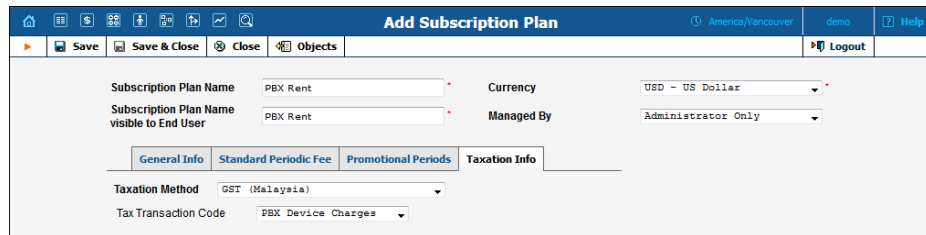
Taxation Method: GST (Malaysia)

Relief Certificate: 123000
Expiration Date: 2015-05-31

Goods:

- Phone Device Charges
- ATA Device Charges
- IAD Device Charges
- PBX Device Charges

- Provides the option to select a special tax transaction code for every subscription and manual transaction, thus ensuring that the required tax rate (6% or 0%) is applied in each case.
- Creates individual tax xDRs for each subscription, manual charge and manual credit.



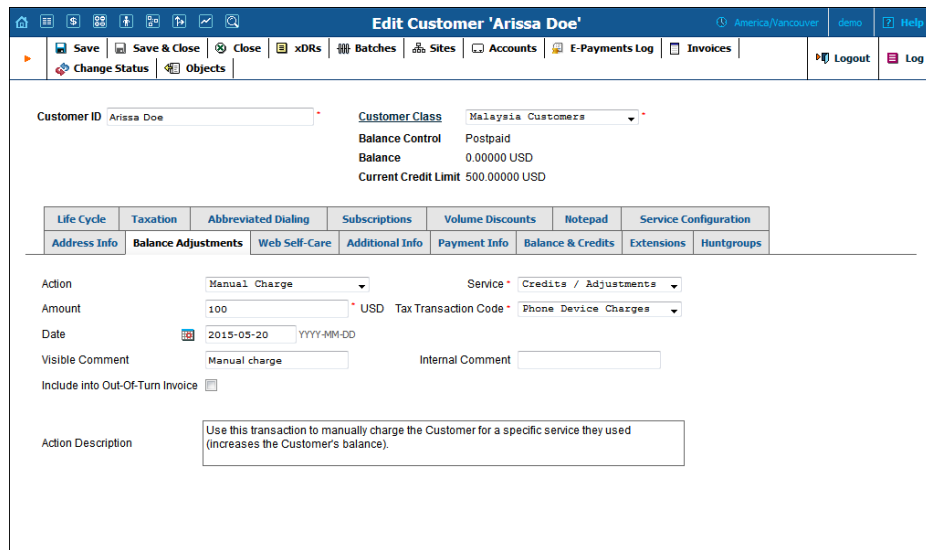
Add Subscription Plan

Subscription Plan Name: PBX Rent | Currency: USD - US Dollar

Subscription Plan Name visible to End User: PBX Rent | Managed By: Administrator Only

Taxation Method: GST (Malaysia)

Tax Transaction Code: PBX Device Charges



Edit Customer 'Arissa Doe'

Customer ID: Arissa Doe | Customer Class: Malaysia Customers

Balance Control: Postpaid
Balance: 0.00000 USD
Current Credit Limit: 500.00000 USD

Life Cycle	Taxation	Abbreviated Dialing	Subscriptions	Volume Discounts	Notepad	Service Configuration	
Address Info	Balance Adjustments	Web Self-Care	Additional Info	Payment Info	Balance & Credits	Extensions	Huntgroups

Action: Manual Charge | Service: Credits / Adjustments

Amount: 100 USD | Tax Transaction Code: Phone Device Charges

Date: 2015-05-20

Visible Comment: Manual charge | Internal Comment:

Include into Out-Of-Turn Invoice:

Action Description: Use this transaction to manually charge the Customer for a specific service they used (increases the Customer's balance).

- 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.).

External taxation of subscriptions

According to tax legislation, taxes are calculated for different types of services individually. If there is a transaction of a defined type (e.g. an international or intra-state call) for a certain service, it is pretty simple.

However, when there is a bundle of services within one product covered by a single subscription (e.g. product “EasyCall” provides 100 free minutes for domestic destinations, 200 minutes of calls to Europe and SIP phone rental for a \$20 monthly fee), the tax calculation is more complicated.

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). Therefore, taxes must be calculated for each service within a bundle as well as separately for each subscription fee.

The performance of this task is handled by the External taxation modules that are integrated with PortaBilling®. It is also the responsibility of the service provider to describe the service bundles (i.e. which services a subscription covers and in what proportion subscription fees are distributed among these services) to the tax company and acquire corresponding tax codes. Then the system will know how the services covered by this subscription should be taxed. Of course, to acquire different taxation codes you can distribute the subscription fees among the services within a bundle in different amounts, and define the activation and cancellation fees for the subscription as well.

After the service description has been provided to the tax company (EZtax® or SureTax®), the company defines the corresponding tax codes for these services and forwards them to the service provider. These codes are then entered into PortaBilling®.

Depending on the taxation company, the tax codes can be of different types. SureTax defines only transaction codes while EZtax divides them into transaction and service codes.

Let's extend the example above and 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.

Additionally, there is a \$10 activation fee that covers equipment costs and a \$5 cancellation fee that covers administrative costs as defined by this subscription.

The selected tax company is EZtax, therefore this information is sent to them.

After the transaction / service codes are acquired from EZtax, an administrator must perform further configurations in PortaBilling® on the **Taxation** tab for a particular subscription:

1. Select the EZtax taxation method from the list.
2. Specify the corresponding EZtax Service / Transaction codes for specific (activation, usual (i.e. a periodic fee of this subscription that is defined on the **Periodic Fee** tab) and cancellation) fees.
3. Save the changes.

Transaction / Service	Value 1	Value 2
Activation	10	15
Cancellation	10	15
Usual	10	15

Once defined, the tax information is stored in the PortaBilling® database. During the tax calculation process it is sent to the tax company where the taxes are calculated and the data is forwarded back to PortaBilling®. It is then added to a customer's invoice, ensuring correct subscription taxation.

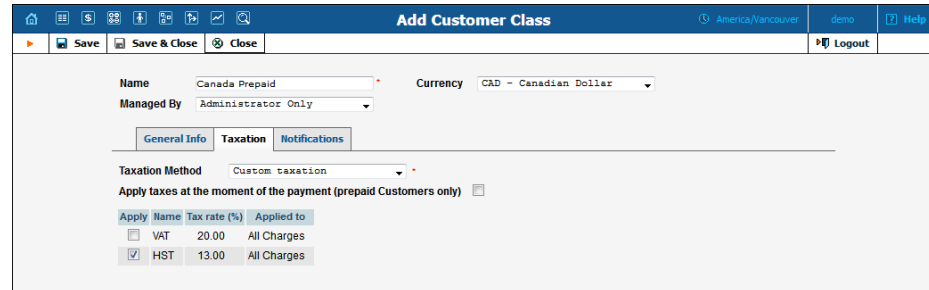
Taxes upon Payment

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).



Apply	Name	Tax rate (%)	Applied to
<input type="checkbox"/>	VAT	20.00	All Charges
<input checked="" type="checkbox"/>	HST	13.00	All Charges

For now, tax calculation upon payment is supported for the Fixed VAT Rate and Custom taxation methods only. In future releases, it will be implemented for other supported taxation plug-ins.

The same tax determination procedure is supported for **Auto-payment**.

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 him / her 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 awhile 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's name 'John Doe (Login: jdoe)'. The main heading is '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

There are buttons for 'Pay Now' and 'Use Other Payment Methods'.

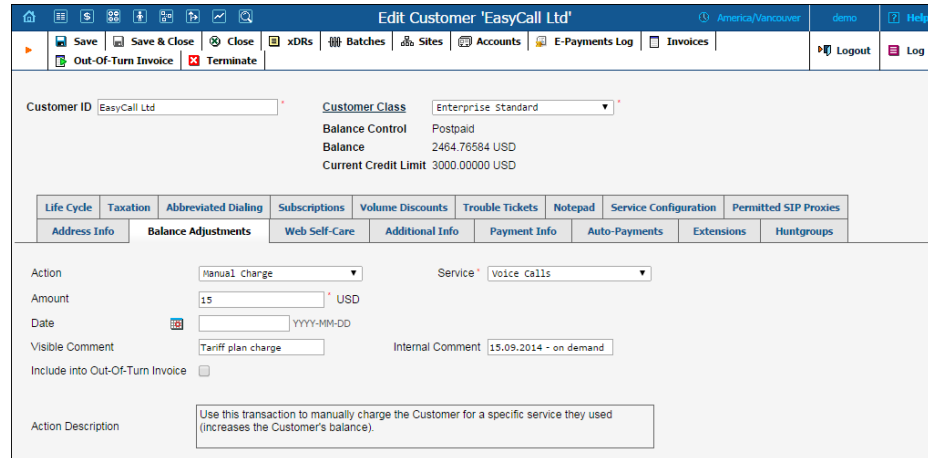
John Doe confirms the payment and his credit card is charged \$11.30. After the payment is processed, two xDRs are recorded in the database: one for the amount paid (\$11.30) and another one for the amount of tax (\$1.30). John's available funds are now \$14.

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.

Manual 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:

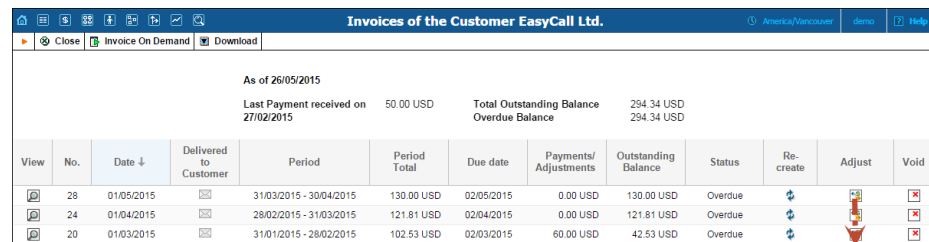
Transaction Type	Description	Effect On Postpaid Customer's Balance	Effect On Prepaid Customer's Available Funds
Manual Charge	Use this transaction to manually charge a customer for a specific service they used.	Increase	Decrease
Manual Credit	Use this transaction to manually give compensation related 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 refund an earlier payment received from the customer (e.g. a check returned by the bank).	Increase	Decrease
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

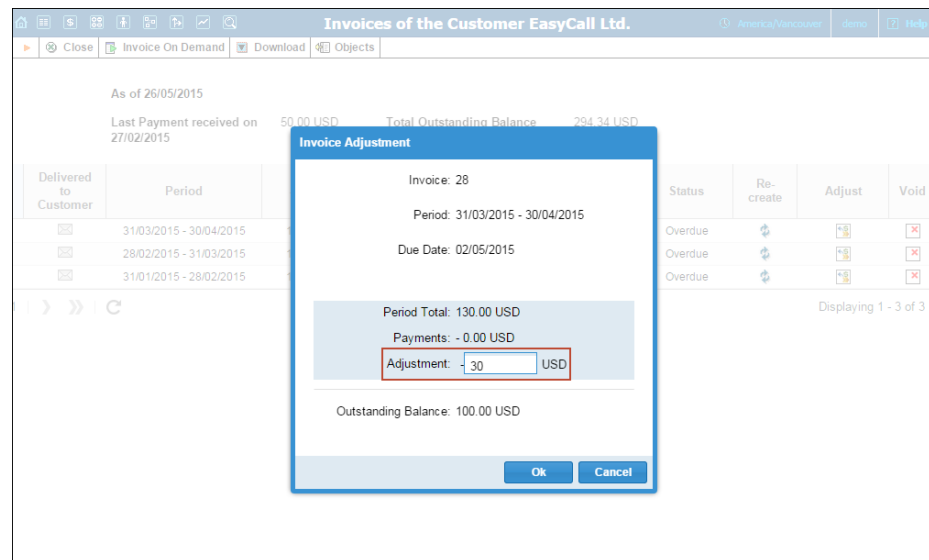
Outstanding Balance Adjustment

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.



View	No.	Date ↓	Delivered to Customer	Period	Period Total	Due date	Payments/ Adjustments	Outstanding Balance	Status	Re-create	Adjust	Void
	28	01/05/2015		31/03/2015 - 30/04/2015	130.00 USD	02/05/2015	0.00 USD	130.00 USD	Overdue			
	24	01/04/2015		28/02/2015 - 31/03/2015	121.81 USD	02/04/2015	0.00 USD	121.81 USD	Overdue			
	20	01/03/2015		31/01/2015 - 28/02/2015	102.53 USD	02/03/2015	60.00 USD	42.53 USD	Overdue			



Invoice Adjustment

Invoice: 28

Period: 31/03/2015 - 30/04/2015

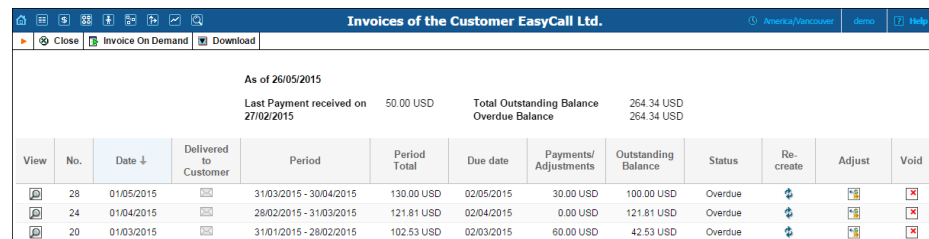
Due Date: 02/05/2015

Period Total: 130.00 USD

Payments: - 0.00 USD

Adjustment: USD

Outstanding Balance: 100.00 USD



View	No.	Date ↓	Delivered to Customer	Period	Period Total	Due date	Payments/ Adjustments	Outstanding Balance	Status	Re-create	Adjust	Void
	28	01/05/2015		31/03/2015 - 30/04/2015	130.00 USD	02/05/2015	30.00 USD	100.00 USD	Overdue			
	24	01/04/2015		28/02/2015 - 31/03/2015	121.81 USD	02/04/2015	0.00 USD	121.81 USD	Overdue			
	20	01/03/2015		31/01/2015 - 28/02/2015	102.53 USD	02/03/2015	60.00 USD	42.53 USD	Overdue			

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 USD is generated and on April 1st, a March invoice for 500 USD is generated. Later in April, the customer disputes some of the charges from his February invoice and is given a waiver for 50 USD. The administrator enters the waived amount and therefore the customer pays 1450 USD. Now his February and March invoices are paid.

These transactions will be reflected in his April invoice:

- Previous Balance – 1500 USD.
- Invoice Adjustment – 50 USD.
- Payment – 1450 USD .

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.

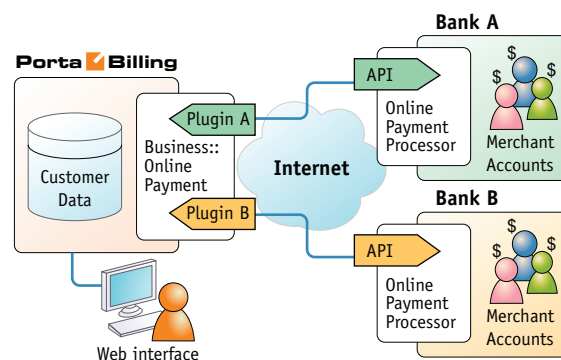
Payments

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, there is 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 about 10 different plug-in modules are supplied with PortaBilling®, and it is easy to create new plug-ins for online payment processors. You can find more information about the Business::OnlinePayment module and available plug-ins at <http://www.cpan.org/modules/by-module/Business/>.

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.

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 and Payment Express 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.

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.

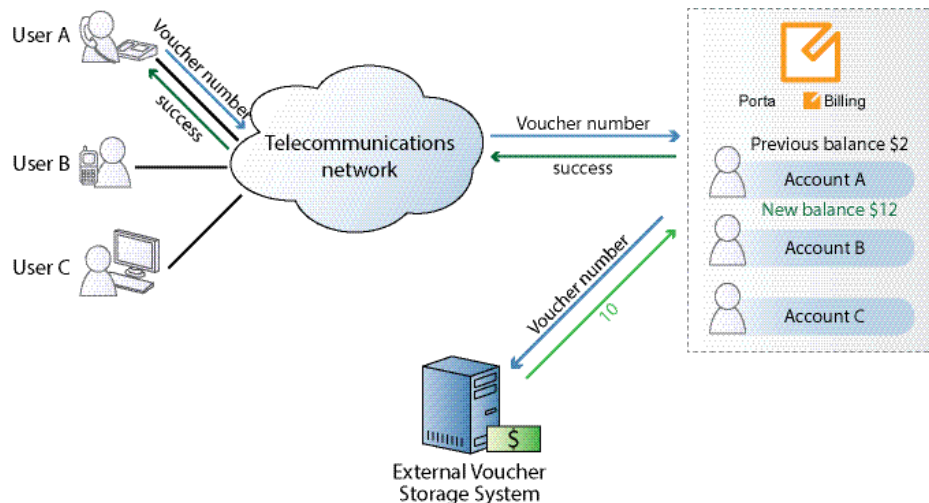
One example is receiving payments made via payment terminals. This payment method is one of the most popular and frequently used in the Russian Federation and other developing countries. The payment process is the following: the customer chooses a service provider on the screen of a payment terminal, 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.

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.



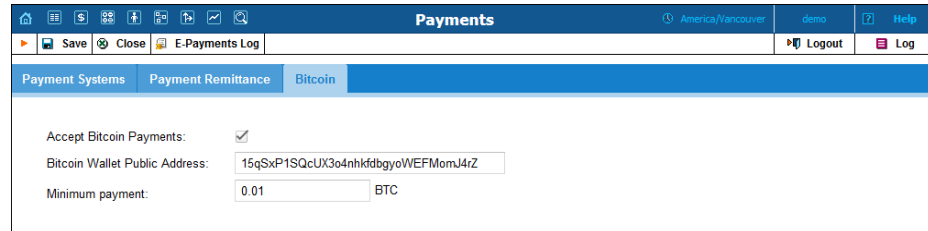
Note that for account top-up with external vouchers the currency of the voucher is assumed to be the same as that of the customer.

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.

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.

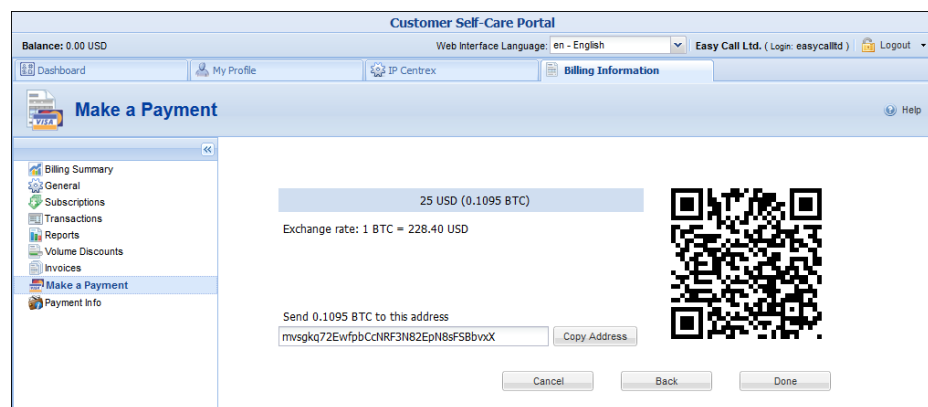


The screenshot shows the 'Payments' configuration page with the 'Bitcoin' tab active. It includes a 'Save' button, 'Close', and 'E-Payments Log' options. The configuration details are as follows:

- Accept Bitcoin Payments:
- Bitcoin Wallet Public Address:
- Minimum payment: BTC

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 screenshot shows the 'Customer Self-Care Portal' with the 'Make a Payment' section. The interface includes a navigation menu on the left and a main content area with the following details:

- Balance: 0.00 USD
- Web Interface Language: en - English
- Easy Call Ltd. (Login: easycallid) Logout
- Navigation: Dashboard, My Profile, IP Centrex, Billing Information
- Section: **Make a Payment**
- Payment Amount: 25 USD (0.1095 BTC)
- Exchange rate: 1 BTC = 228.40 USD
- QR code for scanning
- Send 0.1095 BTC to this address:
- Buttons: Cancel, Back, Done

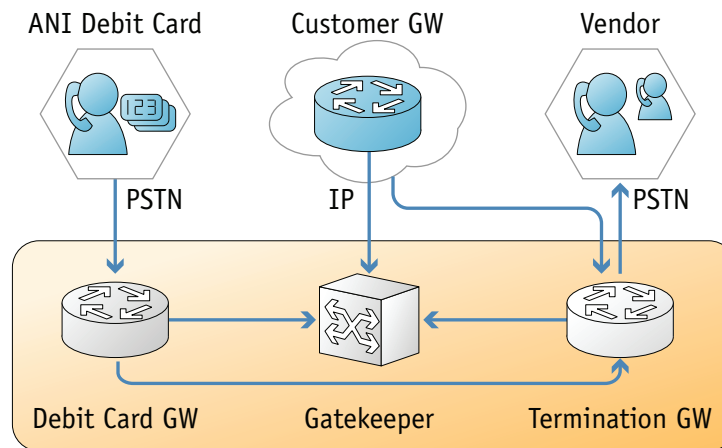
The principal strength of Bitcoin is that it offers anonymity to end users. When a Bitcoin transaction takes place, no personal identifying information such as name and address are required. With Bitcoin support in PortaBilling®, service providers and resellers are able to offer their services to their customers who care about privacy and want it preserved.

4. Voice Services

CDR Mediation

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 CDR record.

Collecting data from the VoIP network



Porta  Billing Owner IP

Gateways in the PortaBilling® owner network are “trusted” gateways. These gateways **MUST** be listed in the Nodes directory and be a RADIUS 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).

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).

Currently, the routing information can be obtained directly from billing by the PortaSIP® Switching Server node (for SIP calls).

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 Algorithm

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 node (Media Server, etc.) which runs that 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 Switching 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 (offline, 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 Switching 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 for IVR Applications

Sometimes there is a need to configure the IVR application routing since some applications require particular PortaSIP® Media Server nodes, or a single node (for example, the *Conferencing* IVR application, which doesn't support multi-node routing).

IVR application routing can be set up in the following ways:

- Automatic – When this type of routing is selected, the system will randomly route a call to one of the PortaSIP® Media Server nodes.
- Manual – When this type of routing is selected, the administrator can set up custom routing for specific PortaSIP® Media Server nodes by defining the order of their use for this IVR application.

Manual routing allows for the configuration of the following scenarios:

- Selecting a single PortaUM® node for where calls made to IVR applications that do not support multi-node routing can be routed.
- Adding PortaUM® nodes to the list and sorting them to route calls sequentially.

For example, if your Media Servers are installed in multiple geographic locations, then it would be desirable to first route calls that arrive from local telcos to the Media Servers located in the same region. This configuration will improve call quality by reducing network delay between the telco's originating gateway and the PortaSIP® Media Server.

- Adding a pool of PortaUM® nodes so that they can receive calls interchangeably.

For example, if you add a pool with two Media Servers to the routing list, then each server has a 50% chance of receiving a call.

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.

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 defines the routes for which categories are available, 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.

Routing algorithm

The routing principle is simple:

- 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; entries with a preference or route category lower than the one in the first huntstop entry will be ignored.
- 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.
- 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 termination 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%).
- 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).

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.

Does PortaSwitch® support LCR?

Yes, PortaSwitch® supports LCR – and much more besides. In fact, “just LCR” is the simplest type of routing PortaSwitch® handles. If you decide not to use routing plans (one default plan for everyone) or routing preferences (same preference for all vendors), then the routing decision will be affected solely by the vendor’s cost. Also, PortaBilling® supports a “profit-guarantee” mode for routing. Only those termination carriers who satisfy your conditions for the minimum required profit (e.g. at least \$0.005 each minute) will be chosen.

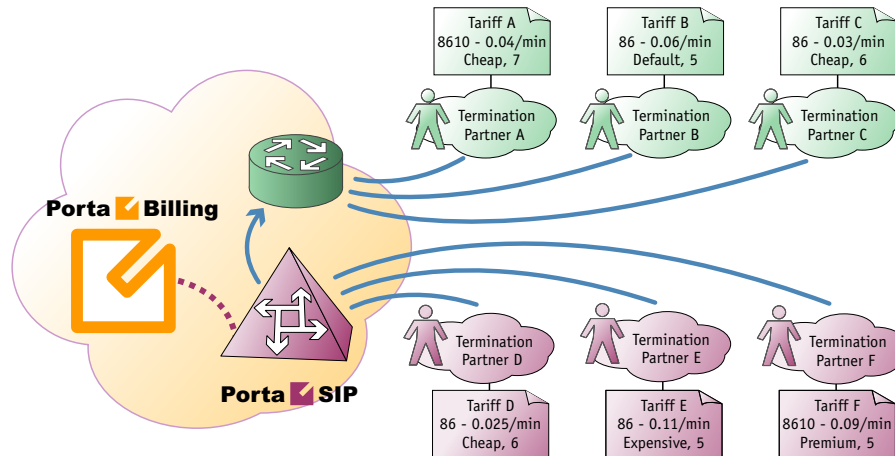
Compare the two illustrations below. The first one shows simple least-cost routing: all of the available carriers are arranged according to cost. In the second illustration, PortaSwitch® routing preferences are used, so that the administrator can re-arrange the routing. In this case, carrier D has been moved down the routing list, while carrier A has been moved up to the top of the list, thus becoming the first route.

#	Route CLID	Route To	Connection	Penalization	Route Category	Preference	Huntstop	Price, EUR	Destination	Country	Description	Tariff	Vendor
1	86	45.32.156.200	Termination to vendor D		Default	5	N	0.02500	86	China	Proper	Vendor D	Vendor D
2	86	76.63.159.25	Termination to vendor C		Default	5	N	0.03000	86	China	Proper	Vendor C	Vendor C
3	8610	76.63.159.25	Termination to vendor A		Default	5	N	0.04000	8610	China	Beijing	Vendor A	Vendor A
4	86	76.63.159.25	Termination to vendor B		Default	5	N	0.06000	86	China	Proper	Vendor B	Vendor B
5	8610	64.67.2.191	Termination to vendor F		Default	5	N	0.09000	8610	China	Beijing	Vendor F	Vendor F
6	86	193.50.123.6	Termination to vendor E		Default	5	N	0.11000	86	China	Proper	Vendor E	Vendor E

Phone Number	Routing Plan	Routing Mode	Apply Penalties	Date and Time
8610234567	All Available Routes	Generic Routing	<input checked="" type="checkbox"/>	YYYY-MM-DD HH24:ML:SS

#	Route CLD	Route To	Connection	Penalization	Route Category	Preference	Huntstop	Price, EUR	Destination	Country	Description	Tariff	Vendor
1	8610	76.63.159.25	Termination to vendor A		Default	6	N	0.04000	8610	China	Beijing	Vendor A	Vendor A
2	86	76.63.159.25	Termination to vendor C		Default	5	N	0.03000	86	China	Proper	Vendor C	Vendor C
3	86	76.63.159.25	Termination to vendor B		Default	5	N	0.06000	86	China	Proper	Vendor B	Vendor B
4	8610	64.67.2.191	Termination to vendor F		Default	5	N	0.09000	8610	China	Beijing	Vendor F	Vendor F
5	86	193.50.123.6	Termination to vendor E		Default	5	N	0.11000	86	China	Proper	Vendor E	Vendor E
6	86	45.12.156.200	Termination to vendor D		Default	2	N	0.02500	86	China	Proper	Vendor D	Vendor D

Routing configuration example



Consider the following example. If you have:

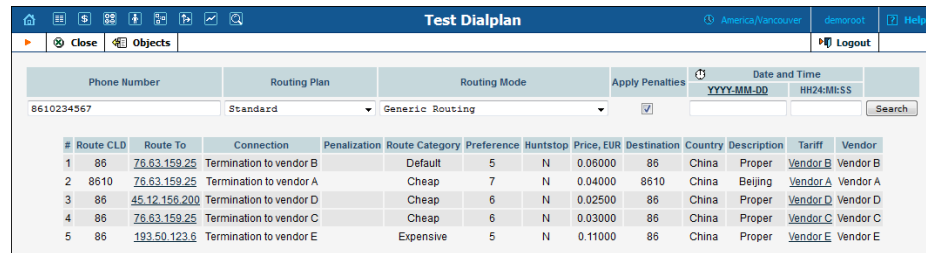
- A “Standard” routing plan, which includes three route categories: “Default” (order 70), “Cheap” (order 40) and “Expensive” (order 10).
- Six vendors (A, B, C, D, E, F) with the following rates (prefix, route category, preference, price):
 - A. 8610, Cheap, 7, 0.04
 - B. 86, Default, 5, 0.06
 - C. 86, 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).



#	Route CLD	Route To	Connection	Penalization	Route Category	Preference	Huntstop	Price, EUR	Destination	Country	Description	Tariff	Vendor
1	86	76.63.159.25	Termination to vendor B		Default	5	N	0.06000	86	China	Proper	Vendor B	Vendor B
2	8610	76.63.159.25	Termination to vendor A		Cheap	7	N	0.04000	8610	China	Beijing	Vendor A	Vendor A
3	86	45.32.156.200	Termination to vendor D		Cheap	6	N	0.02500	86	China	Proper	Vendor D	Vendor D
4	86	76.63.159.25	Termination to vendor C		Cheap	6	N	0.03000	86	China	Proper	Vendor C	Vendor C
5	86	193.50.123.6	Termination to vendor E		Expensive	5	N	0.11000	86	China	Proper	Vendor E	Vendor E

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 go further down the routing list (via other carriers; Vendor B in our example) to 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 our 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 billing of these calls.

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® now 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).

Attribute	Value
Description	Termination to X-Telecom
Type	Calls to Vendor via SIP
Service Type	Voice Calls
Routing Criteria	High priority
Active	<input checked="" type="checkbox"/>
Capacity	100
Limit Utilization by Capacity	<input checked="" type="checkbox"/>
Start Utilization Balancing After	75 calls

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. Capacity	Load Threshold	Current Load	Handicap	Adjusted 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 Plan Selection

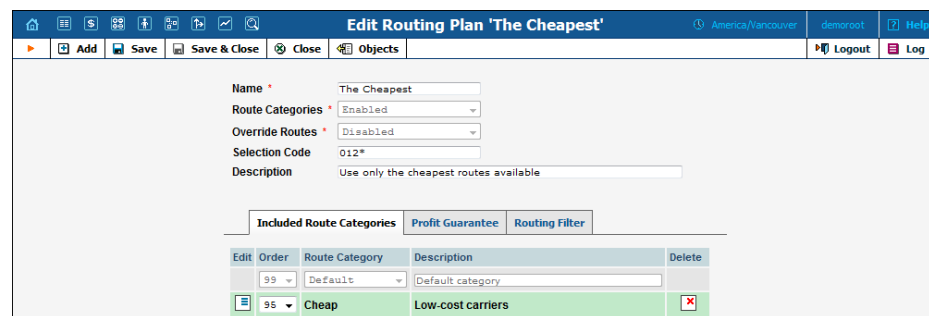
There are two possibilities for choosing which routing plan is used to find available vendors and terminate a call:

- Each account is assigned a default 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.



Edit	Order	Route Category	Description	Delete
	99	Default	Default category	
	95	Cheap	Low-cost carriers	X

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 it 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.

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.

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 Successful 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.

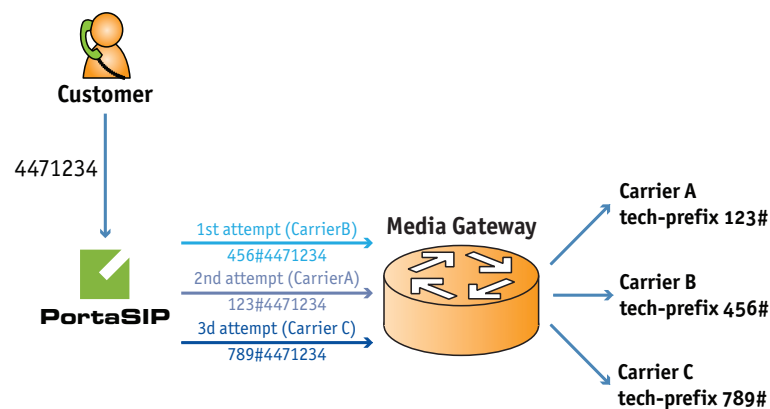
-  RED – The route was manually unblocked; this status will remain until the next time interval for computing statistics.

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 IPIPGW, or any of the SBCs available on the market as a SIP-H323 convertor. 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:

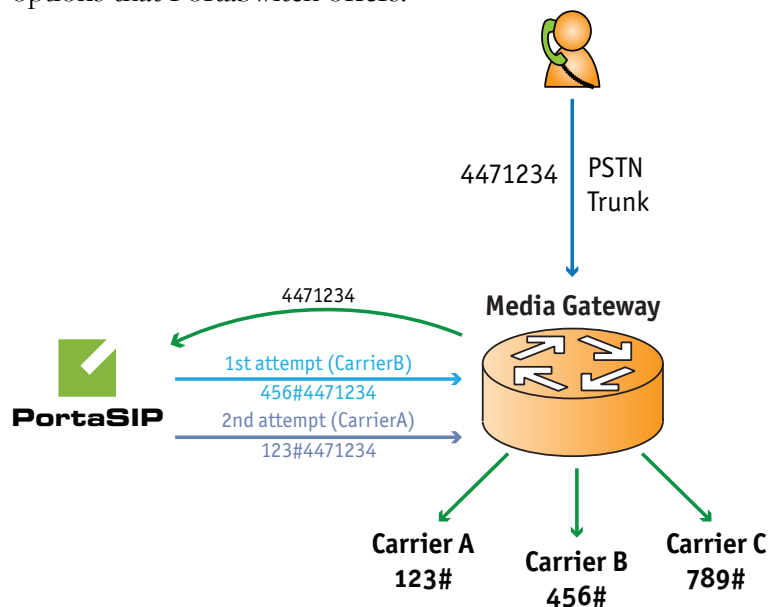
```
Route 1: 4412345678 @ <IP address of carrier A>
Route 2: 2#4412345678 @ 192.168.0.3
Route 3: 4412345678 @ <IP address of carrier C>
Route 4: 4#4412345678 @ 192.168.0.3
```

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

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. When a customer makes a phone call, it arrives from outside your network and the translation rule specific to that customer 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 translation 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 translation rule is applied, so that the number becomes 4202123456. This is the number that will be stored in the xDR information.

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® Web Reference Guide](#). 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).

Your network numbering plan

The key to avoiding problems with number formats is to choose a certain number format as the standard for your network and make sure that calls travel on your network only in this format. The ideal candidate for such a format is E.164 (of course it is highly recommended that you use this same format in billing as well).

Customer-based translation 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. Customer-based translation rules allow you to convert a number from a format specific to a particular customer. Customers can even manage their dialing format themselves via the self-care interface.

Customer-based translation 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, you can convert it to 42021234567 and find the correct rates for the 420 prefix.
- This number will be returned to the node which requested it.

Connection-based outgoing translation rules

If your vendor requires a special number format (e.g. tech-prefix), it is possible to apply this rule to convert the number. When billing returns a list of routes, the phone numbers for routes for this connection will be converted. This only works for a routing model in which the VoIP node (e.g. PortaSIP®) requests billing for routing information. If your gateway uses dial-peers or an external gatekeeper for routing, then you must configure number translation there.

Connection-based translation 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. 7834#42021234567). Now it is necessary to convert this number to the proper format for billing (4202134567), which may be done using connection translation 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).

Connection translation rules are also used on **VoIP from Vendor** 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).

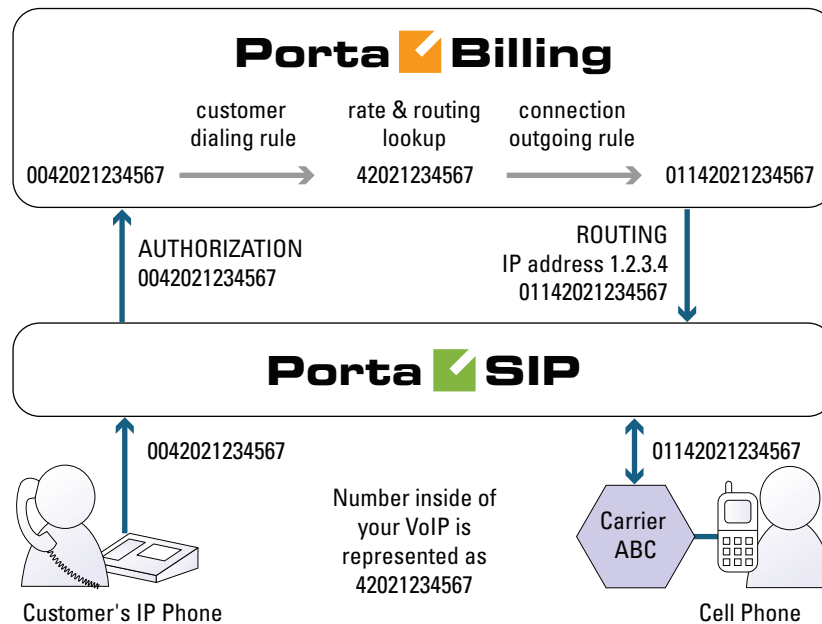
Node-based translation rules

These serve the purpose of converting a number from a custom format used by the customer into billing's internal format during authorization, depending on the gateway. For example, on a gateway in Prague, Czech Republic, there may be the translation rule “strip leading 00”, while on a gateway in Moscow, Russia, the rule will be “strip leading 810 or replace leading 8 with 7”.

Currently, node-based translation rules are only kept for backward compatibility with older releases, since they have been made obsolete by customer translation rules. They will most likely be removed from future releases.

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.

DID Provisioning and Billing

Number Management

The DID inventory in PortaBilling® is extended to allow ownership of a particular DID number by a specific customer. This will facilitate working with the customer's list of DID numbers, even if they are not currently assigned to any accounts. It will also allow end-customers to perform these operations 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 ITSPs 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 allocating 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 owner 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.

Billing for DID Numbers Usage

Every DID number belongs to an owner 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 him / her.

This allows you to set up different pricing for your customers 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” owner 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 bi-weekly 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” owner 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

When defining the price to charge the end user for a DID as a part of the DID owner information, the following parameters are available:

- **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.
- **Base Periodic Fee** is a configurable fixed amount that is always included in the recurring (monthly) charge.
- **Percentage Periodic 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.

As a result, the total charge for the DID would be:

- **At activation** – DID provider activation charge + Activation fee.
- **Monthly** – DID provider monthly charge + the markup, consisting of the Base periodic fee and the Percentage periodic

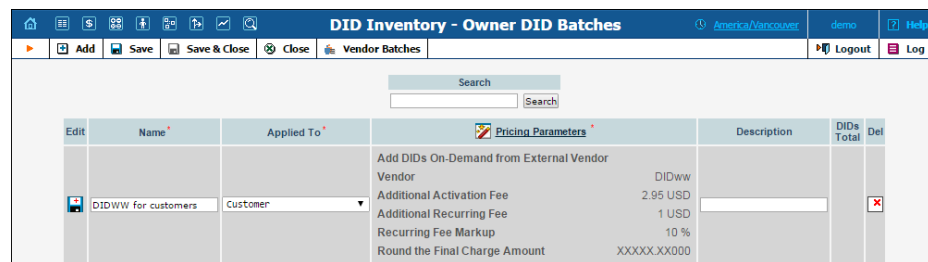
markup. This formula is likely to result in some “inconvenient” numbers, such as 3.289, which are usually rounded up.

Rounding

In order to obtain “marketable” amounts to be presented on the customer’s invoice it is possible to specify the rounding pattern for the charged amounts. This will specify the last “pre-set” digits in the amount.

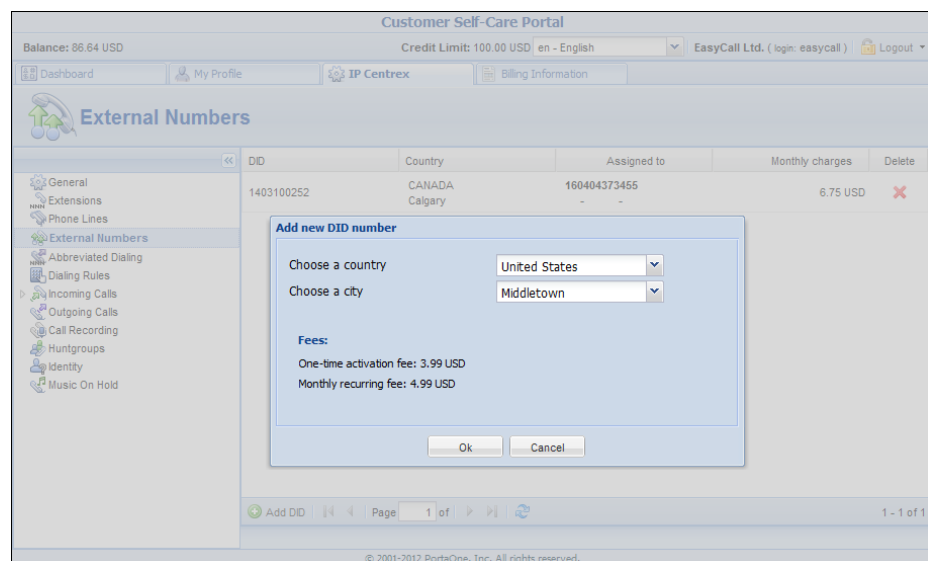
For instance, in a situation where number 12065551234 (which has a \$1.00 activation fee and a \$3.00 monthly fee charged by the DID provider) is provisioned by an end user, the DID pricing setup would be as follows:

- \$2.95 for activation markup,
- \$1 for base periodic markup,
- 10% for percentage periodic markup,
- and X.99 for rounding (this parameter produces a charge amount with .99 at the end).



Edit	Name*	Applied To*	Pricing Parameters*	Description	DIDs Total	Del
	DIDWW for customers	Customer	Add DIDs On-Demand from External Vendor Vendor: DIDww Additional Activation Fee: 2.95 USD Additional Recurring Fee: 1 USD Recurring Fee Markup: 10 % Round the Final Charge Amount: XXXXX.XX000			X

The customer will then be charged \$3.99 ($\$1.00 + \$2.95 = \3.95, rounded up to \$3.99) at activation and \$4.99 ($\$3.00 + \$3.00 * 10\% + \$1.00 = \$3.00 + \$0.30 + \$1.00 = \4.30, rounded up to \$4.99) each month.



DD	Country	Assigned to	Monthly charges	Delete
1403100252	CANADA Calgary	160404373455	6.75 USD	X

Add new DID number

Choose a country:

Choose a city:

Fees:

One-time activation fee: 3.99 USD

Monthly recurring fee: 4.99 USD

Ok Cancel

In the case of a customer having 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 him / her. 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 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 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 an owner 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.

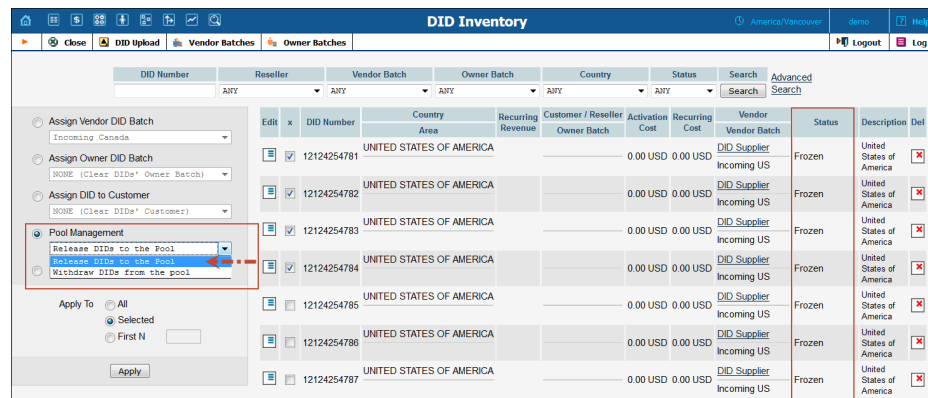
When a customer signs up for a product directly or via a reseller, no costs for DID provisioning apply to him or her.

Thus, with completely free of charge DIDs you perform simplified management of internal DID numbers.

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 new **Frozen** status. When frozen, DIDs are displayed on the administrator’s DID inventory page only.



DID Number	Reseller	Vendor Batch	Owner Batch	Country	Status	Description	Del
12124254781				UNITED STATES OF AMERICA	Frozen	United States of America	X
12124254782				UNITED STATES OF AMERICA	Frozen	United States of America	X
12124254783				UNITED STATES OF AMERICA	Frozen	United States of America	X
12124254784				UNITED STATES OF AMERICA	Frozen	United States of America	X
12124254785				UNITED STATES OF AMERICA	Frozen	United States of America	X
12124254786				UNITED STATES OF AMERICA	Frozen	United States of America	X
12124254787				UNITED STATES OF AMERICA	Frozen	United States of America	X

To allow resellers to provision DID numbers, the administrator simply releases them to the pool of available DIDs.

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 owner batches, and
- not assigned to customers or resellers.



In order to retrieve DID numbers from the pool, resellers must have owner batches with a DID markup assigned to them, since they are charged according to the pricing parameters defined within these owner batches.

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.

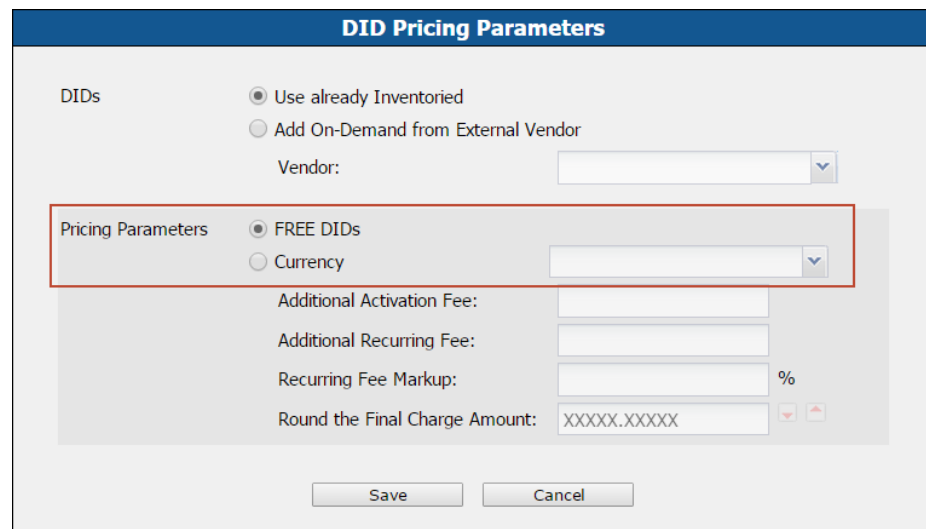
When the product is released, you unfreeze the DID numbers by assigning them to respective owner batches.

Or, if you want your resellers to provision DIDs along with the new product, release some of these DIDs to the DID pool.

Thus, this frozen DID Number functionality and the DID pool make DID provisioning more flexible and effective.

Assigning DID Numbers to Owner Batches

You can provision DID numbers to your customers either with your markup or free of charge. The type of owner batch a DID number is assigned to is what determines whether it is free or not. The owner batch type is selected during the batch creation in the **Pricing Parameters** dialog box.



DID Pricing Parameters

DIDs

Use already Inventoried

Add On-Demand from External Vendor

Vendor:

Pricing Parameters

FREE DIDs

Currency

Additional Activation Fee:

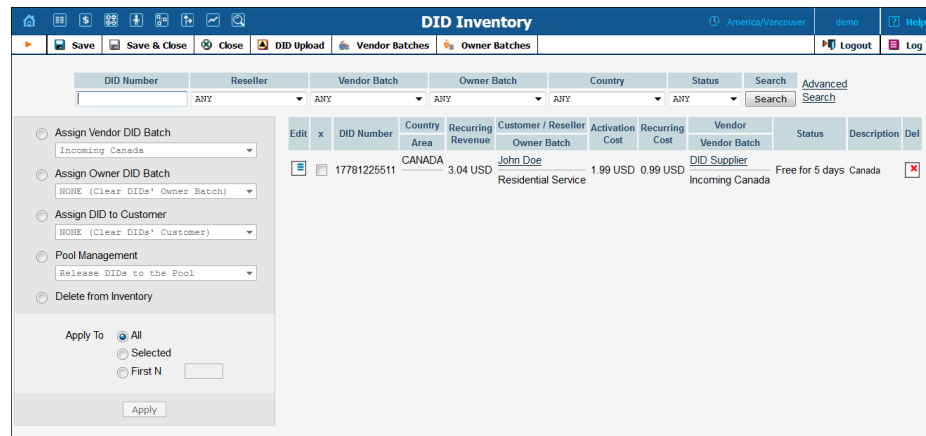
Additional Recurring Fee:

Recurring Fee Markup: %

Round the Final Charge Amount: XXXXX.XXXXX

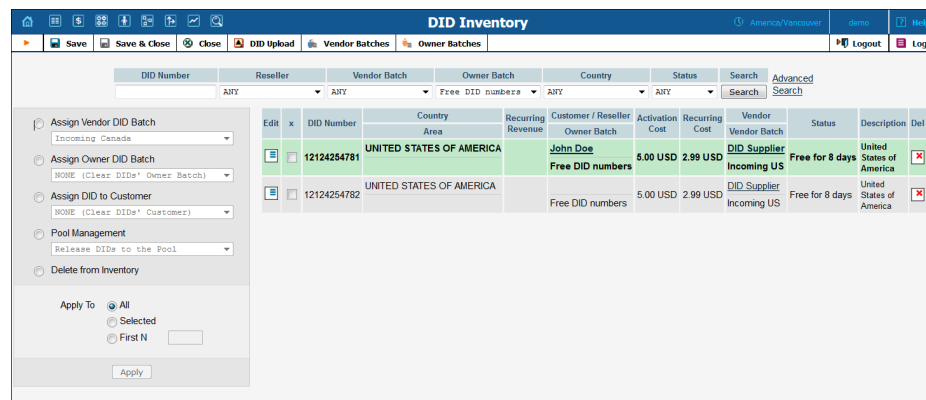


To properly manage DID numbers and control your costs and revenue, it is important to assign each DID number to an appropriate owner batch prior to allocating this number to a customer.



When a DID number is assigned to a customer from an owner batch with a DID markup, the markup is applied to the customer.

Owner batches reflect a DID markup specific for either an administrator or a reseller. An administrator does not see any owner batches created by resellers because those contain resellers’ markups.



When added to a free owner 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 owner batch is assigned to them (e.g. DID number 16045552879 has a \$10 activation cost associated with it and is assigned to the free owner 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 owner batches they are currently assigned to to other owner batches. If a DID number is not in use, it can be reassigned to any owner 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 owner batch (i.e. to free owner batches or to owner batches with DID markup).

Provisioning DIDs to Customers

Depending upon your business model, you can provision DIDs 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 owner 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 owner 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 an owner batch to a customer and if this customer has another owner batch specifically assigned to him, the charges defined in the customer's owner batch will be ignored. Instead, the charges defined in the owner batch that a DID number belongs to will apply.

If customers provision DID numbers themselves (i.e. by using on-demand DID provisioning), it is required that the owner batch having the DID markup be assigned to them because they are billed according to the pricing parameters defined in this specific owner batch.

A DID number assigned to a customer and used as an account or account alias remains assigned to this customer after the alias or the account is terminated, and the customer is still charged for this number usage.

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 [Handbooks Collection](#) for detailed instructions about different ways of DID number provisioning to customers.

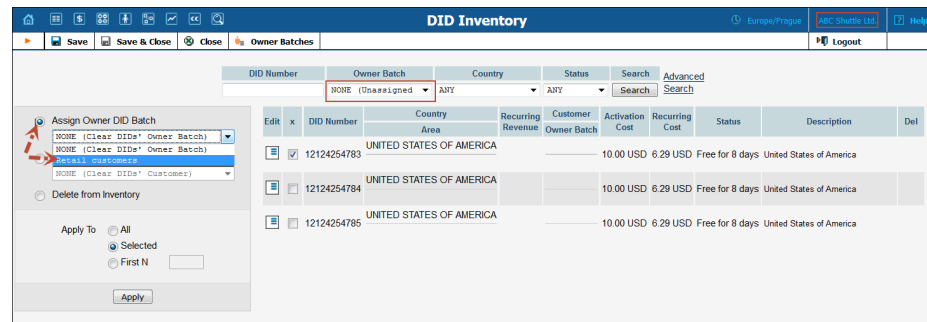
Reseller DID Provisioning

Access to the DID pool allows resellers to use as many DID numbers as they need and then provision these numbers to their customers.

Consider the following example:

A reseller ABC Shuttle Ltd. needs 20 DIDs for a corporate customer. Instead of asking the administrator for these numbers and waiting for their assignment, the reseller does the following:

1. Goes to the DID inventory page and creates the owner DID batch;
2. Selects **Unassigned DIDs** from the **Owner Batch** list on the search bar;
3. Selects the country and the amount of required DIDs;
4. Assigns them to the owner DID batch and then to the required customer.



To release a DID number back to the DID pool (e.g. a reseller has accidentally assigned a wrong DID number to the owner batch), the reseller un-assigns this number from his owner batch. Thus this DID number becomes available for other resellers.


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 PortaSwitch® side will reserve and allocate that number on the DID provider side and provision it inside of PortaSwitch® (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 allows us 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>) provider – and more providers will be added soon based on requests from our customers.

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 C. 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.
4. 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 the PortaBilling® system 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 **Number Porting** button on the **Account Management** page. In the request dialog window two parameters must be mentioned:

- The phone number that the customer wants to use;

- The date by which this request must be completed.

The screenshot displays the 'Account Info / Retail Customer 'John Doe'' interface. At the top, there are navigation buttons like 'Save', 'Save & Close', 'Close', 'xDRs', 'E-Payments Log', 'Terminate', 'Number Porting', 'Objects', 'Logout', and 'Log'. Below this, the account details are shown: Account ID 160404373456, Balance Control dropdown, Blocked checkbox, and Balance 0.00000 USD. A grid of tabs is visible, including Life Cycle, Subscriptions, Volume Discounts, Notepad, Service Configuration, Forward, and Phone Book. The 'Account Info' tab is active, showing customer details like Name (John Doe), Type (Credit), Service Password (jymtshz2), E-Mail, Customer Site (None), and Batch (None). A 'Create Request' dialog box is overlaid on the page, with fields for 'PortIn Number' (3125705500) and 'Desired Due Date' (2015-06-10 00:00). The dialog has 'Create Request' and 'Cancel' buttons.

PortaBilling® sends the porting request to the number porting company and waits for confirmation. The number porting company 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** tab) and to his system (**Port-Out** tab) by clicking the **Number Porting** button on the **Account Info** page from the main menu.

Number Porting					
Port-In	Port-Out				
	Account ID	Porting Number	Desired Due Date	Status	Cancel Request
	2015678324	201-706-5879	2015-04-26 11:12:00	PROCESSING	✘
	2015678324	201-704-4892	2015-04-16 12:31:00	PROCESSING	✘
	2015678424	201-696-0158	2015-04-16 12:31:00	PROCESSING	✘
	2016778524	209-400-0478	2015-05-26 10:02:00	PROCESSING	✘
	2036678424	209-397-7895	2015-04-10 08:08:00	PROCESSING	✘



NOTE: The **Number Porting** button is only available when **NumberPorting** is enabled on the Configuration Server. To use Neustar for number portability, **NumberPortability_Neustar** must also be enabled.

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.

Premium Number Billing

This is actually quite similar to the situation with ported numbers described above. PortaBilling® has to perform a lookup (either in a local database or an external one) to determine what the routing destination for the number is and which rating code should be used to charge your customer.

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.

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.

Bandwidth-aware Limitation of Allowed Concurrent Calls

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/Quintum 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.

IP Device Provisioning and Inventory

If you provide your VoIP customers with IP phone equipment, 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

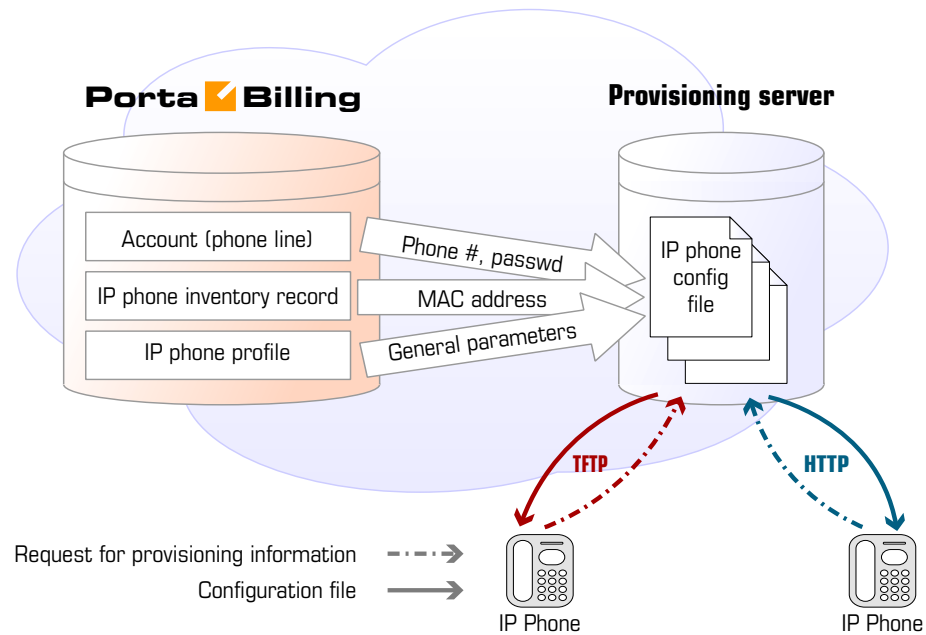
This approach is a fundamentally different one. Instead of attempting to contact an IP phone and change its parameters (pop method), the initiative is transferred to the IP phone itself. The device will periodically go to the provisioning server and fetch its configuration file.

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 IP phone 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 IP phone 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.



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.

IP Phone Inventory

The IP phone directory 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 IP phone inventory.

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.

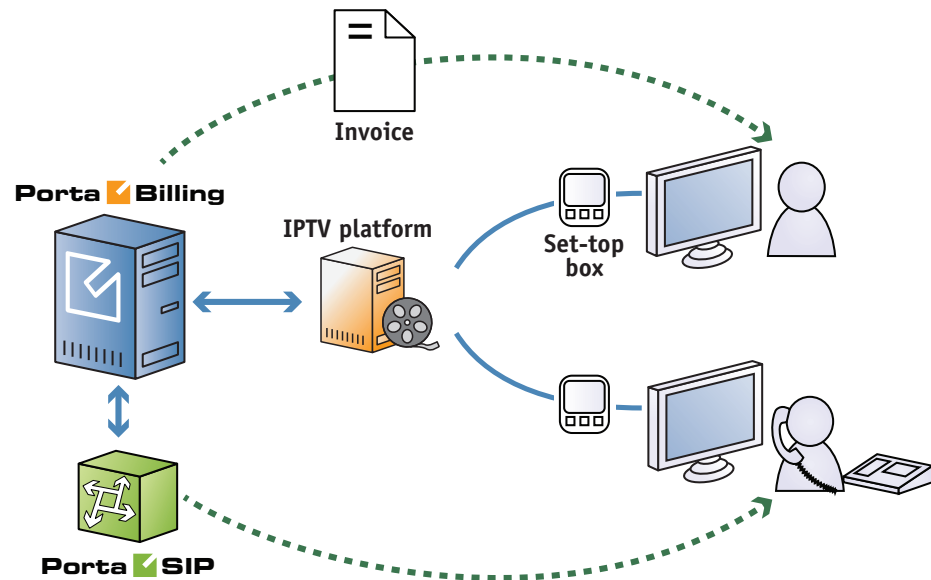
NOTE: From the billing point of view, a SIP server is not different from any other gateway - it also issues authentication or authorization requests, and sends accounting information after the call is completed.

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.

5. Converged Services

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.

Currently the Matrix IPTV and Minerva platforms are supported.

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.

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 use Internet access services.



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.

- 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.

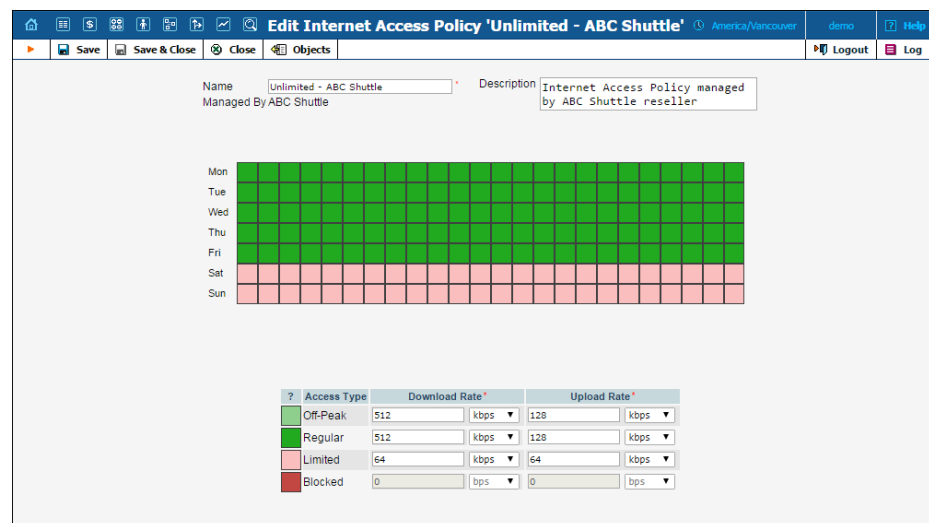
Controlling Bandwidth

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.



Name: Unlimited - ABC Shuttle
 Description: Internet Access Policy managed by ABC Shuttle reseller

Day	Access Type
Mon	Off-Peak
Tue	Off-Peak
Wed	Off-Peak
Thu	Off-Peak
Fri	Off-Peak
Sat	Limited
Sun	Blocked

Access Type	Download Rate*	Upload Rate*
Off-Peak	512 kbps	128 kbps
Regular	512 kbps	128 kbps
Limited	64 kbps	64 kbps
Blocked	0 bps	0 bps

There are four available types of Internet access:

- Normal access (default).

- Access during off-peak hours (often referred to as "turbo" access).
- 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:

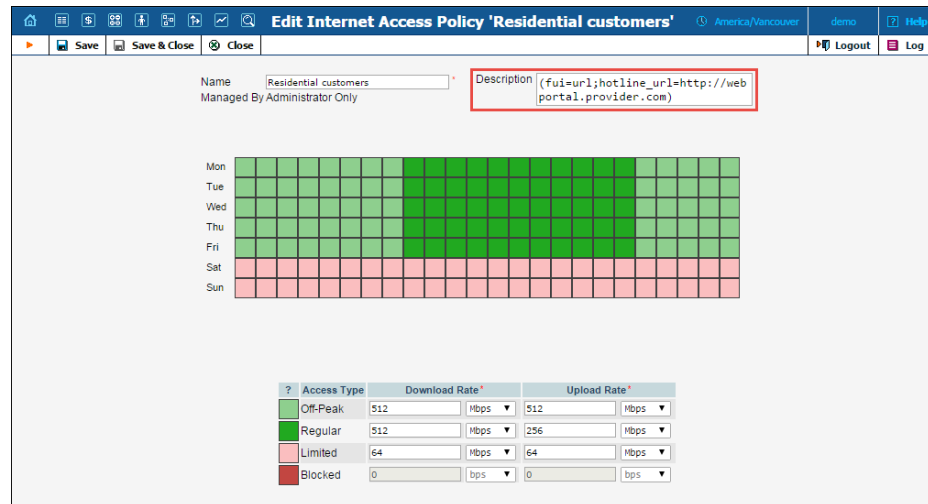
- A scheduling table, which designates what type of access is used during each time period.
- Specific values for the allowed upload / download speed for each type.

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).

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 either assign it to a product or directly to an account. Input the main configuration parameters in brackets in the **Description** field:



Find all parameters in [Appendix F](#) of this document.

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 PortaBilling 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.

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

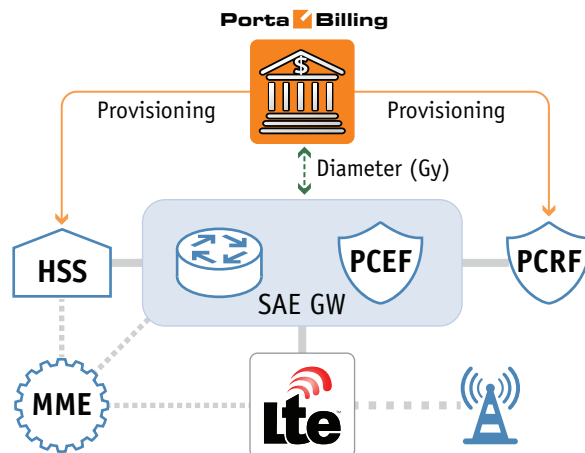
In cases when a customer needs to purchase an entire IP-network from the ISP, PortaBilling® provides the option for configuring the routing information for the customer's IP-networks on the NAS. To do this, the administrator must indicate the network prefix, the length of the network prefix mask and the gateway IP address in the **Routed Networks**

attribute. This routing information will then be automatically propagated to NAS each time the customer connects.

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.

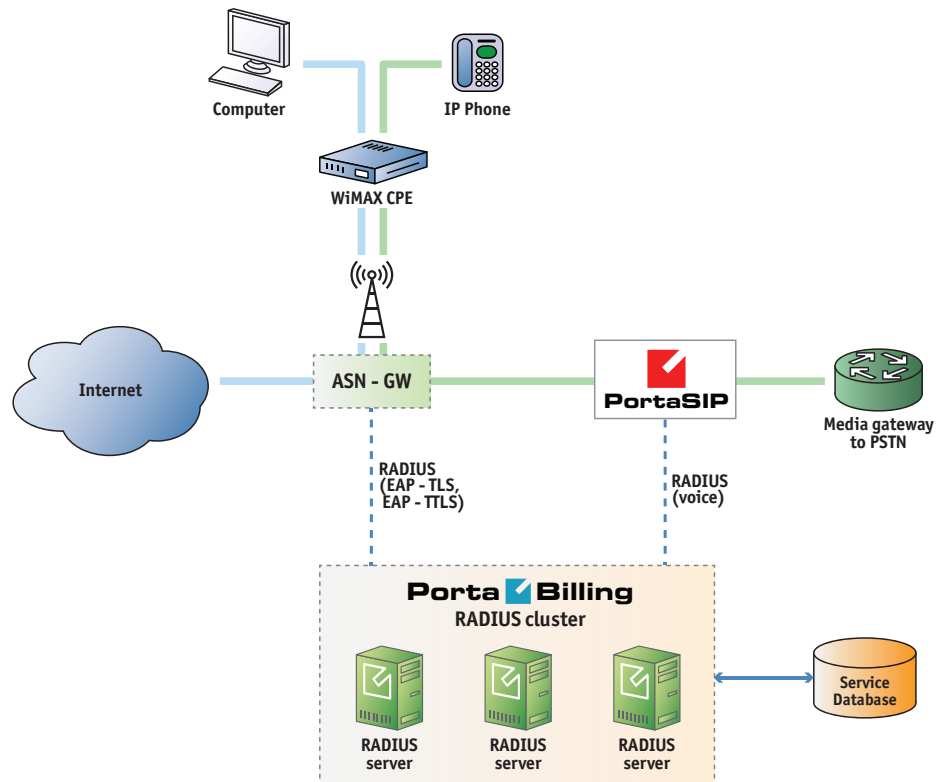


In terms of the network, PortaBilling performs the function of an OCS (Online Charging System). It communicates with SAE-GW via Diameter (Gy) for real-time session authorization and rating. The same Gy interface is used to disconnect a customer exceeding his credit limit or switch him to a different class of services when the usage quota is used up.

PortaBilling can also provision customer data into HSS (to allow centralized user authentication across the network) and PCRF (to allocate the correct policy/QoS parameters to the customer).

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.

Integration with Bridgewater Service Controller

PortaBilling® is integrated with Bridgewater, a service controller for WiMAX networks.

WiMAX services with Bridgewater

PortaBilling supports the following billing models for WiMAX via Bridgewater:

- Preset bandwidth with unlimited data transfer amount, fixed recurrent fee.
- Preset data transfer quota, fixed recurrent fee.

In both cases, the user connects to the ASN gateway, which sends requests to and receives authorization from Bridgewater Service Controller. Bridgewater uses account information from its DB, which in turn, is pre-populated and kept current by PortaBilling.

Account Provisioning

PortaBilling can automatically provision WiMAX accounts into Bridgewater. It means that:

- An account created in PortaBilling will be automatically created in Bridgewater with a Profile² assigned that corresponds to the PortaBilling Product.

² Profiles are to be predefined on Bridgewater side as part of the setup process.

- When an account is suspended due to non-payment or reaches a credit limit in PortaBilling, the account is automatically blocked in Bridgewater.
- When an account is unsuspended or balance returns below a credit limit, the account is automatically unblocked.
- When an account is terminated in PortaBilling, it gets removed from Bridgewater.

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 can serve as the account management and billing platform for such an MVNO.

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

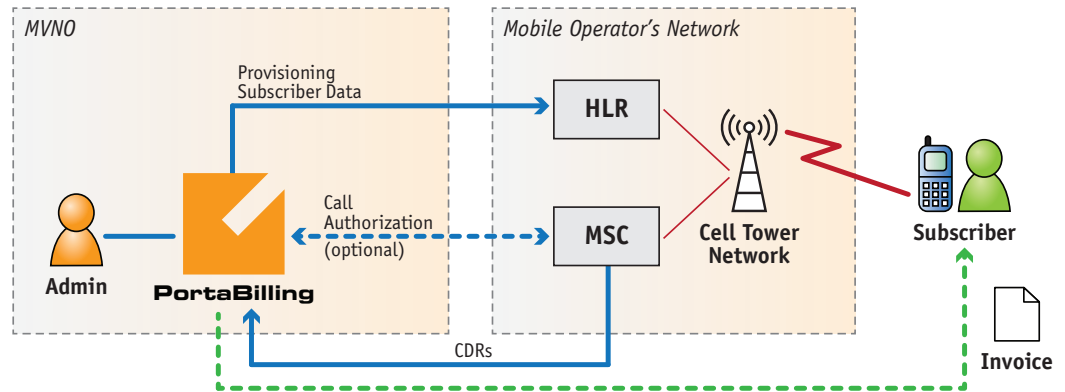
When an account/customer is created by the MVNO, it should be provisioned to the HLR (Home Location Register) of the host mobile operator's system. Account blocking/unblocking, deletion, and so on should also be mapped to the mobile operator.

This is achieved by using External System Provisioning Framework (Sokoban) 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.

Also, in some cases the Mobile Switching Center (MSC) can send authorization requests to PortaBilling for *every* separate call in order to provide balance control. In this case, before establishing a call the MSC will send a call authorization request (typically using the RADIUS protocol) to PortaBilling, and, based on the reply from PortaBilling, it will either allow or deny the call. Real-time call authorization is an important requirement for providing prepaid mobile services as an MVNO.

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 CDR import is then performed.

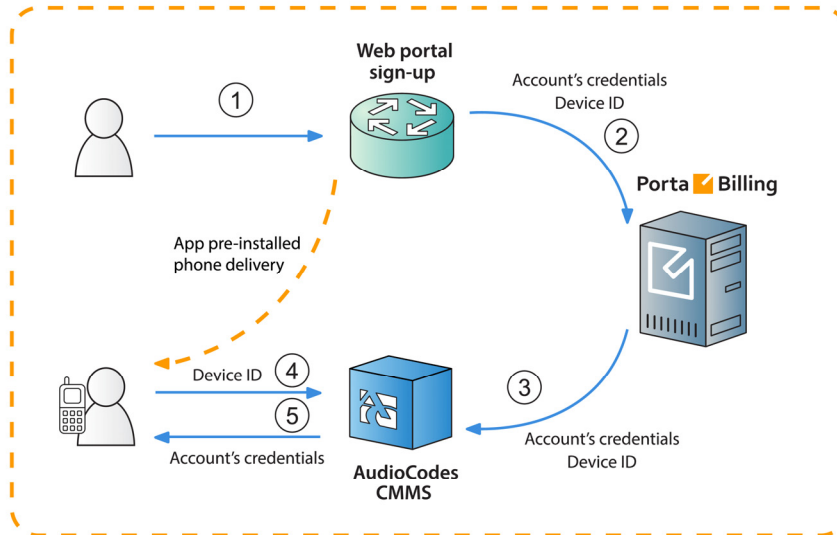


Billing customers and sending out invoices is done directly by PortaBilling, without involving the host mobile operator.

Integration with AudioCodes Mobile Application

PortaBilling® can provision the AudioCodes mobile application for service providers and their resellers. Auto-provisioning makes the configuration of services quick and convenient for end users while saving time and expense for service providers.

This multifeatured AudioCodes mobile application provides a wide range of services: voice and video calls, SMS, IM and more. To provision the application, PortaBilling® is in constant interaction with the AudioCodes CMMS (Clients' Management Server) which is responsible for customized profiles, user level control, time-based licensing and other functions. CMMS, in turn, is in constant interaction with users' mobile phones and tablets, so when a user's information is changed on the PortaBilling® side (e.g. user's credentials), it is quickly updated on the user's side, too.



MVNOs

The AudioCodes mobile application registration and configuration process for the MVNO business model is detailed in the following example:

John Doe wants to subscribe to EasyCall's services. **(1)** He visits the service provider's website and makes a request for a new mobile phone and a list of services to be provided with that phone (e.g. call forwarding, voicemail, etc.). **(2)** A new account is created for John Doe in PortaBilling® that contains information such as account ID, password, IMEI (device ID), etc. **(3)** At the next stage, PortaBilling creates the account in the AudioCodes CMMS with all of the required parameters – the account's credentials, IMEI, etc. The phone with the installed AudioCodes mobile application is then sent to the user. John Doe receives his new phone and switches on the AudioCodes application. **(4)** It automatically connects to the CMMS, provides the IMEI (device ID) and requests the user's credentials. **(5)** The AudioCodes CMMS parses this information and provides those credentials. Thus, the mobile application is configured, John Doe has avoided having to manually set it up and the requested services can be used immediately.

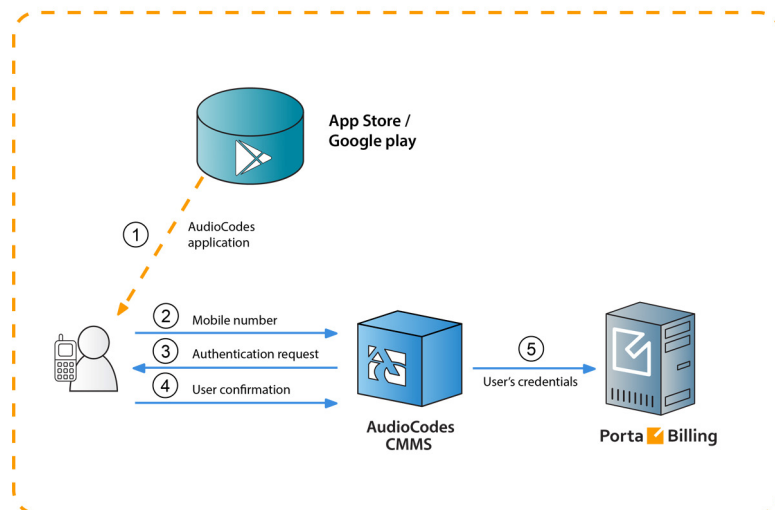
NOTE: AudioCodes CMMS is compatible with Android and iOS only.

OTT service providers

The business model for downloading the AudioCodes mobile application from the App Store or from Google play is applicable to a wide circle of service providers. For this model, creation of a new account in PortaBilling® is initiated by registration of a new user in AudioCodes mobile application. The whole process is described in this following example:

The TelePlus service provider will provide voice calls, messaging and other custom services via their own SIP mobile application. TelePlus obtains the mobile application from AudioCodes and brands it as TelePlus.

(1) John Doe downloads the TelePlus application. The first time he runs it, the application requests his mobile number. (2) He inputs his number and the application sends it to the CMMS server. (3) The server generates a special security code and texts it to John Doe via SMS to confirm that the owner of that number indeed made the request. (4) John Doe enters this security code which is then passed to CMMS. (5) The server accepts the code and sends the request to PortaBilling® with the user's credentials. Then a new account is created in PortaBilling® with John Doe's mobile number used as the account ID. Starting from this moment, PortaBilling® charges John Doe for service usage.



The account information configuration and the updates for the AudioCodes mobile application occur immediately and automatically. Auto-provisioning decreases the workload on the service provider and ensures continuous and smooth service provision, leading to user satisfaction with service quality.

Integration with Cisco PDSN

PortaBilling® can be used by CDMA operators to provision, bill and monitor mobile broadband services controlled by Cisco PDSN (Packet Data Serving Node for CDMA networks). To do this, Cisco PDSN connects to PortaBilling® via RADIUS.

Integration and operation flow

The integration of PortaBilling® with Cisco PDSN is similar to its integration with NAS – the principles of integration are almost the same. Though PDSN is used only for CDMA networks, the main difference between services provided via PDSN and a “traditional” NAS is the user’s mobility. Users may not be stationary and thus switching from one base station to another while the original Internet session is connected. PortaBilling® ensures that even if there are various individual sub-sessions, a single xDR is produced at the end.

Provisioning

To set Cisco PDSN for PortaBilling® you need to carry out settings on the configuration server and on the administrative web interface (tariffs, nodes, etc.)

Integration with CISCO SCE Service Control Engine

PortaBilling® is integrated with Cisco SCE, a network element specifically designed for carrier-grade deployments that require high-capacity stateful application and session-based classification and control of application-level IP traffic, per subscriber.

This solution facilitates content-based charging and allows the ISP to granularly track Internet usage, permitting varied pricing schemes for truly differentiated IP service offerings, while at the same time keeping costs low.


Consider a few examples of services that you could offer using this solution:

- Offer high speed Internet service to your business customers but disallow bandwidth-consuming peer-to-peer (P2P) applications.
- Provide residential customers with standard Internet speed, but offer enhanced latency-sensitive services such as online gaming, by assigning a high priority for identified users of this service.

PortaBilling® currently works with the Cisco SCE integrated with PDSN nodes that already allow CDMA operators to raise their mobile broadband services to new heights. In the future, this solution can be enhanced to support various equipment bundled with Cisco SCEs, thereby benefitting ISPs that provide other Internet services.

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 rate³ 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.

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.”

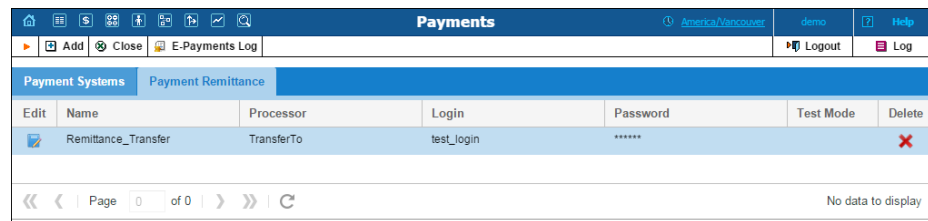
In order to address all of these challenges, the PortaSwitch administrator configures a special “mapping” tariff that flexibly defines the amounts allowed to be added to a remote user’s balance and the corresponding amount to be debited from a PortaSwitch account. For instance, in order

³ At the time of writing

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 39 USD, so taking into consideration exchange rate fluctuations, remittance processor markups and his own profit, the administrator will set the value at US\$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



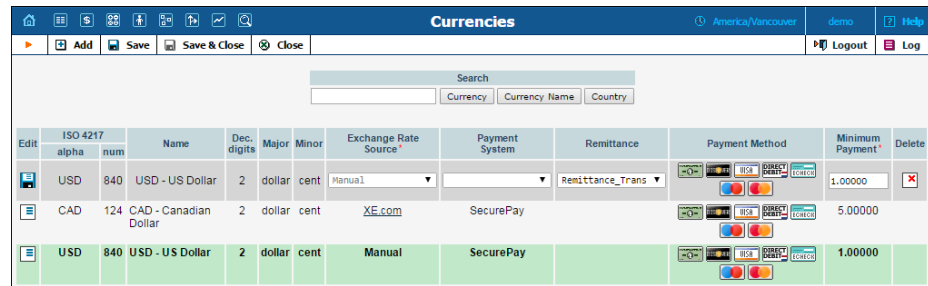
Payments						
Add Close E-Payments Log					Logout Log	
Payment Systems		Payment Remittance				
Edit	Name	Processor	Login	Password	Test Mode	Delete
	Remittance_Transfer	TransferTo	test_login	*****		

Navigation: Page 0 of 0. No data to display.

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.



Edit	ISO 4217 alpha num	Name	Dec. digits	Major	Minor	Exchange Rate Source*	Payment System	Remittance	Payment Method	Minimum Payment*	Delete
	USD 840	USD - US Dollar	2	dollar	cent	Manual		Remittance_Trans		1.00000	
	CAD 124	CAD - Canadian Dollar	2	dollar	cent	XE.com	SecurePay			5.00000	
	USD 840	USD - US Dollar	2	dollar	cent	Manual	SecurePay			1.00000	

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: TransferToCCC (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.

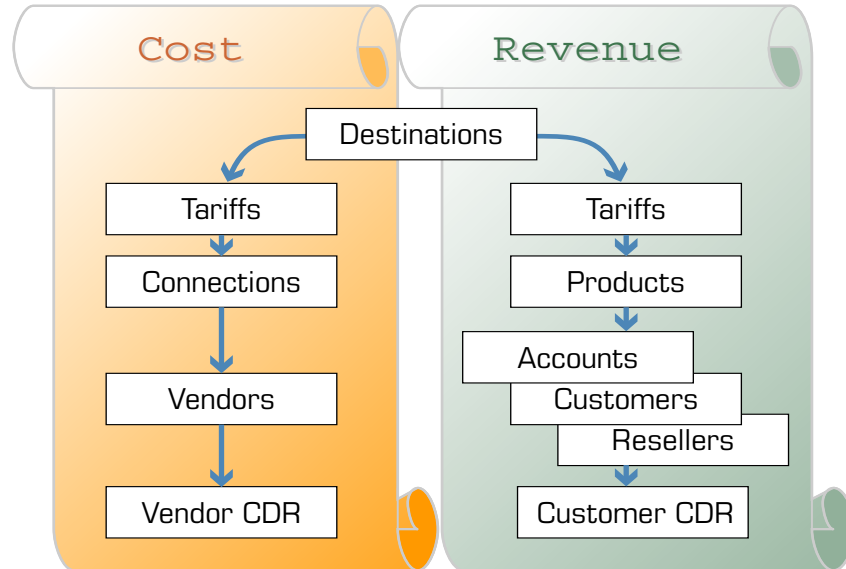
Using the payment remittance system

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.

6 ■ Entering Data in PortaBilling®

Cost / Revenue

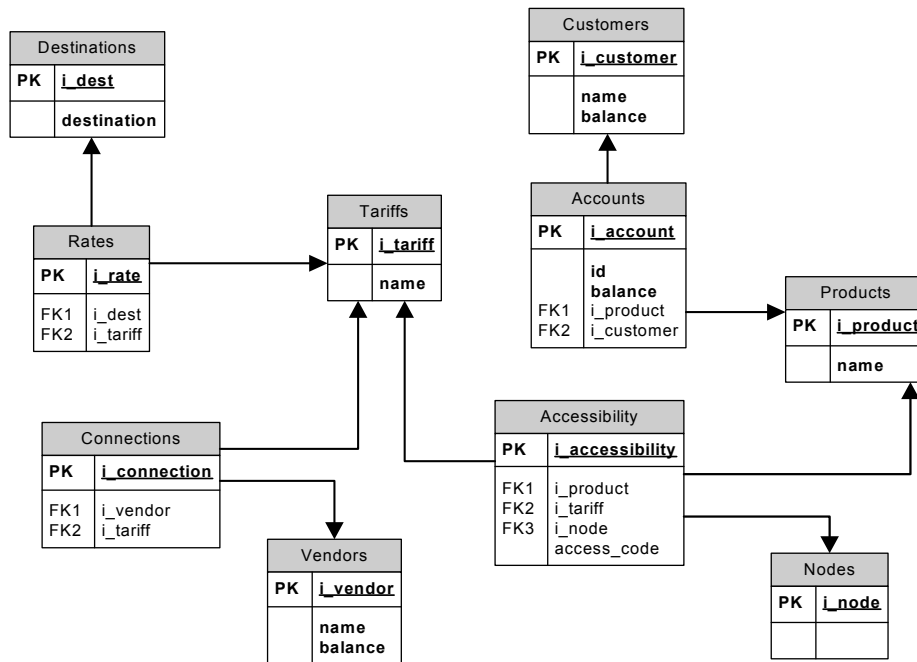
There are two independent parts of the billing configuration:



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.

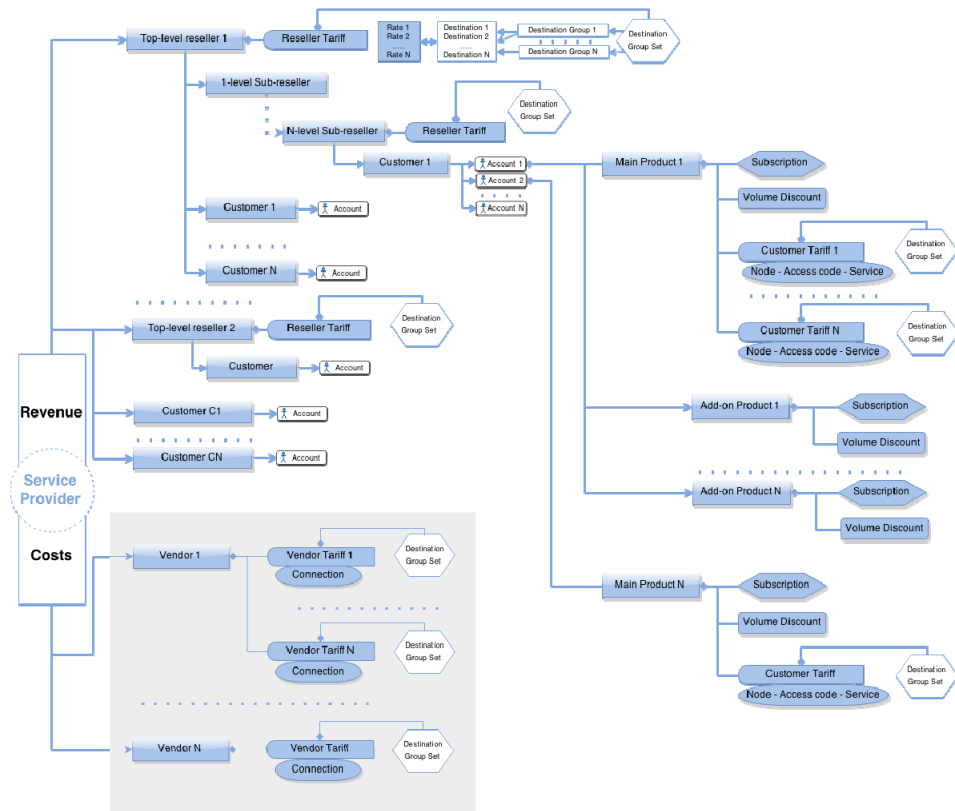
PortaBilling® Data Model Concept

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.

7 ■ Management and Monitoring Tools

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.

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).

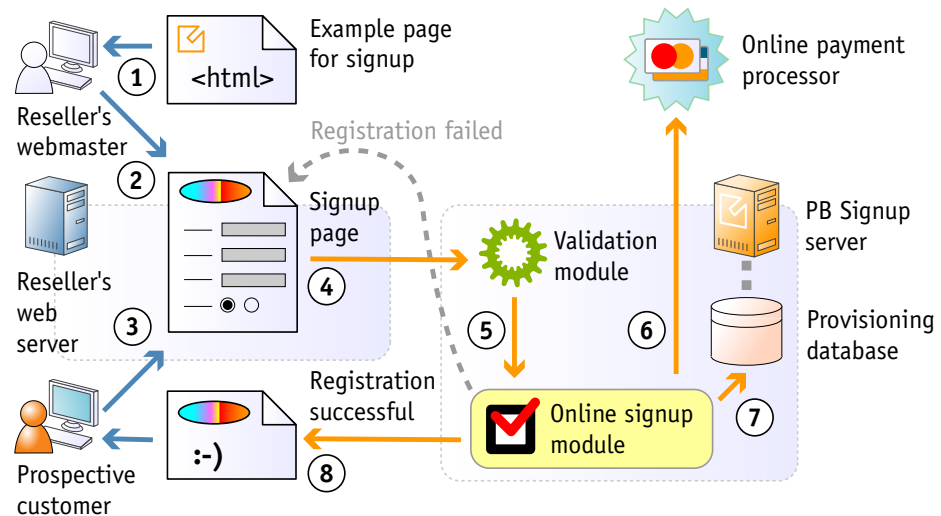
Date and Time Formats

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.

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 payment, 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.

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 *How to ...* 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.

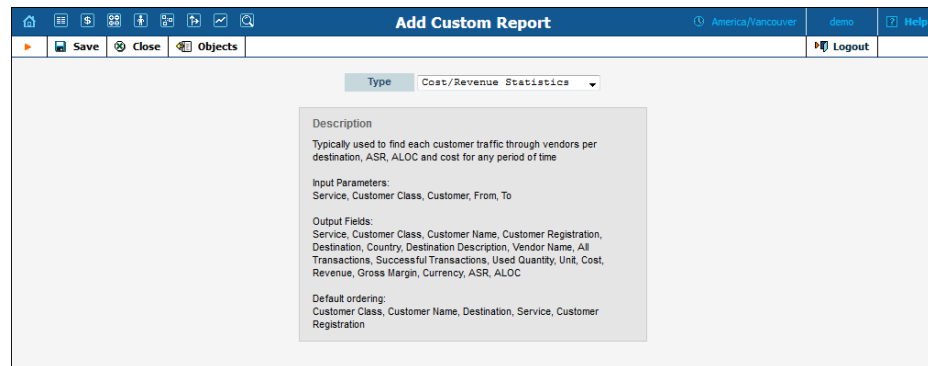
- You may schedule a query to be executed at a certain moment in time or periodically, with the report results delivered to you via email. This allows you to automate reports that you use on a continuous basis. For instance, you may schedule a “Total minutes per customer” report for 3 a.m. every day; 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 invoices (per customer and per periods of 0-30, 31-60, 61-90 and >90 days).
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 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.
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).
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 Routing Info – LCR Blending in the Unified PortaSwitch Handbook Collection .
Traffic Usage Trending Report	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.

Trouble Ticketing System

In order to meet customers’ expectations for quality of service, it is important to keep detailed information about problems which they currently have or had in the past. In addition to information about the service configuration and rating (stored in PortaBilling), you must also keep track of issues reported by a particular customer and the full history of communication and comments on each issue. This allows you to re-use information obtained earlier about the problem (one of the most frustrating things for customers is being asked the same questions all over

again), receive important analytical and reporting data, detect patterns of how problems emerge, and re-use troubleshooting methods successfully applied in the past.

This function is performed by a trouble ticketing system, namely, Best Practical's RT. This is a robust enterprise-grade ticketing system used by many companies worldwide (PortaOne, Inc is one of them). You can find more details about the architecture, functionality and features of RT on Best Practical's [website](#).

RT provides the ability to keep information about individual incidents (tickets):

- Each ticket is associated with a customer and has a unique ID (ticket number) and status (e.g. new, open, resolved). There are some other ticket properties (e.g. the administrator to whom this ticket is currently assigned), but their use generally depends on how the internal ITSP and customer service processes are organized.
- Each ticket has one or more requestors assigned to it. A requestor is a person who communicates with a helpdesk representative. When the helpdesk representative posts new information to the ticket via a **reply**, it is received by all requestors for that ticket via email. (Customer service may also exchange **comments**. These are for internal use only, and are visible only to helpdesk staff, not to requestors.)
- Incoming email communications from administrators or requestors are automatically recorded to the corresponding ticket and then re-sent to the other parties subscribed to that ticket. It is possible for administrators to review the full ticket history from the web interface, while requestors can view the status of tickets and their contents from the web interface (obviously, they can only see **reply** messages). The web interface can also be used to post new messages to the tickets.


PortaBilling is fully integrated with RT, making it a component part of the overall CRM solution. This means that although RT is installed on a separate server and maintains its own database, there is no need to duplicate the entry of any information or log in into each system separately. Information about administrators (helpdesk) and customers is automatically propagated to RT, while information about a particular customer's tickets is dynamically retrieved from RT and presented on the PortaBilling web interface. Thus information about an existing ticket may be retrieved, or a new ticket created, with a single mouse click.

PortaBilling Server Health Monitoring

Entire System Load

These statistics are accessible from the main admin menu via the **System load** link in the **Statistics** section.

This graph shows the number of calls registered by a production system every 15 minutes: normal calls with duration >0 (green areas) and calls with zero duration (red line). The most recent information appearing on the right-hand side of the graph is an hour old or less.



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. (See picture below, call types 3 and 4.)

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.



Main and Replica MySQL Server Statistics

These statistics are accessible from the main admin menu via the **System load** link in the **Statistics** section. Click the  **DB Load** icon to see the information.



These two graphs show values for the status of the MySQL servers. The green line indicates the number of queries processed by the server every 15 minutes, while the blue one shows the number of threads running on the server.

Queries

The graph indicates how many times clients have queried the database. On recommended hardware, this value may exceed 200 without any difficulties. High peaks with a number of requests over 2,000 could indicate improper configuration of the system or temporary problems. Also, these peaks can appear when replication has been restored after a system fault.

Threads

Normal values for these graphs are:

- for the Main database server: ≤ 2 .
- for the Replica database server: ≥ 1 and ≤ 3 .


Normally, these values will depend on the processes running on the server, such as the Radius daemon, the Apache httpd server, statistics collection tools, and so on.

Larger values can indicate other client connections to the server. If the number of threads on the replica server is 0, this means the replication process is down and requires administrator attention.


VoIP Network Performance Statistics

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 **Statistics** section in the main menu.



After you select a vendor and a particular connection, you will see its load graph:



This graph shows the node load, setup time and ASR (Average Success Rate) statistics. The green area indicates node load, the red shows relative setup time, and the blue line indicates ASR. Information appears on this graph 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

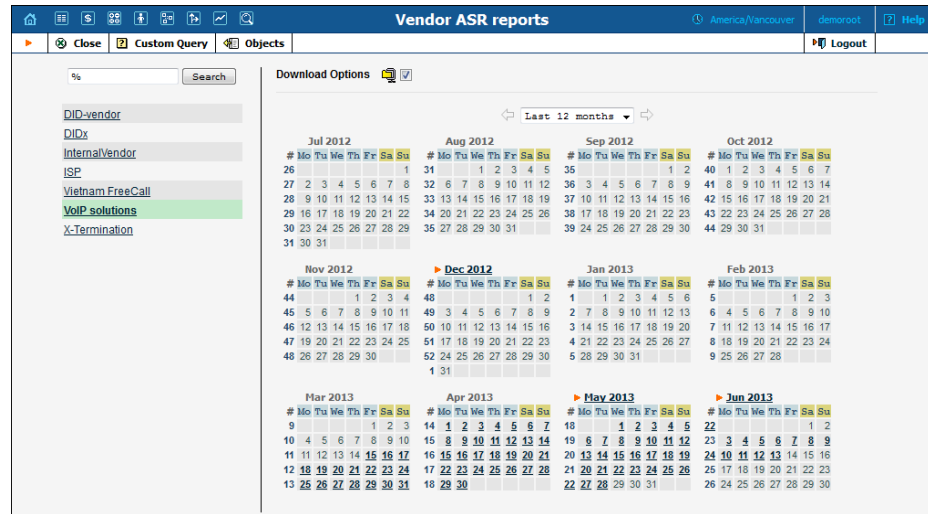
Assumption: Connection load less than 5% is not representative for ASR calculating.

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 ($\text{capacity} * 900 \text{ sec}$).

ASR Statistics



You can download pre-calculated statistics in .csv format (for previous days) or obtain data online using **Custom Query**. For pre-calculated statistics, click on the vendor at left, then choose the statistics period from the calendar. Pre-calculated statistics will look like the following:

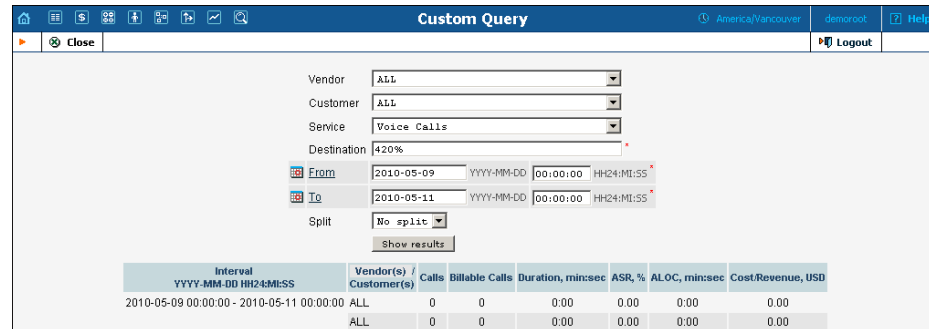
	A	B	C	D	E	F	G	H	I
	Destination	Country	Description	Calls	Billable Calls	ASR, %	Duration, min:sec	ALOC, min:sec	
1	1		North America	587	264	44.97445	1449:49	5:29	
2	SUBTOTAL			587	264	44.97445	1449:49	5:29	
3	420	CZECH REPUBLIC	Proper	590	297	50.33898	1722:49	5:48	
4	SUBTOTAL			590	297	50.33898	1722:49	5:48	
5	380	UKRAINE	Proper	1171	626	53.45858	3554:22	5:40	
6	SUBTOTAL			1171	626	53.45858	3554:22	5:40	
7	TOTAL			2348	1187	50.55366	6727:0	5:40	

The following data is available:

- Total number of calls.
- Number of calls with non-zero duration (billable calls).
- ASR.
- Total call duration.
- Average Length of Call (ALOC).

Numbers are aggregated per destination prefix, with a subtotal per country and a total for all calls.

Should you need statistics for today, or a report with certain other parameters (for example, divided by hour, so that you see how the ASR evolved during the day), you can use the **Custom Query** report. This extracts data directly from the database, so use it with caution in order not to overload the server.



Interval YYYY-MM-DD HH24:MI:SS	Vendor(s) / Customer(s)	Calls	Billable Calls	Duration, min:sec	ASR, %	ALOC, min:sec	Cost/Revenue, USD
2010-05-09 00:00:00 - 2010-05-11 00:00:00	ALL	0	0	0:00	0.00	0:00	0.00
	ALL	0	0	0:00	0.00	0:00	0.00

Note that both ASR and Cost/Revenue parameters are shown on the **Custom Query** report.

Billing Statistics

Customer xDRs

Lists of all customers' xDRs are calculated daily, and are split into files according to each customer's billing period. These pre-calculated statistics can be automatically mailed to the customer. They are also available for download by the customer on the self-care pages, and for your staff on the admin interface.

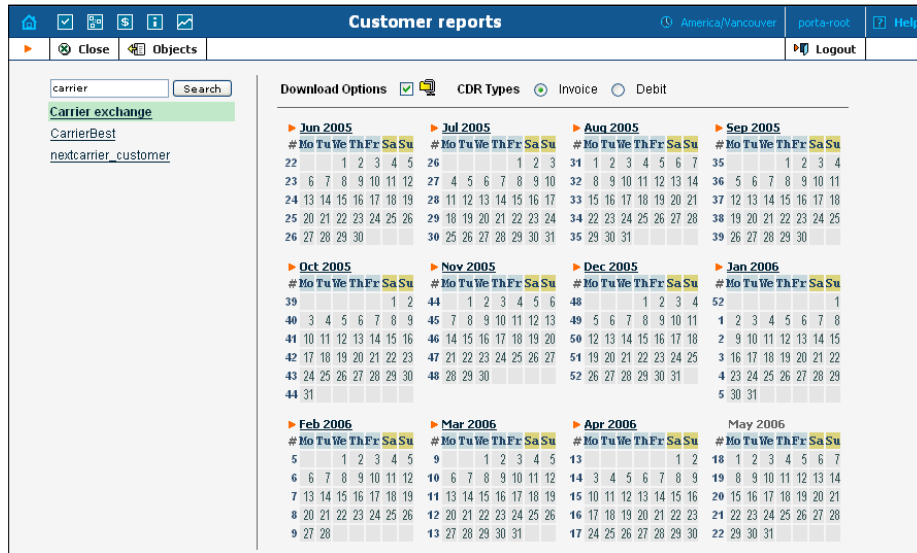
Different types of xDR files are available, depending on the type of customer:

xDR files for retail customers

- **Invoice** – All activities (calls, data access, messaging, etc.) on credit accounts of this customer. These transactions will be included on the customer's invoice for the corresponding billing period.
- **Debit** – Transactions related to activities on debit accounts of this customer. Since debit accounts are prepaid, charges applied to them do not affect the customer's balance or invoice. Therefore, they are included in a separate file so that the customer can easily monitor activities on his debit accounts.

xDR files for resellers

xDRs calculated using the reseller’s wholesale tariff. These are the charges applied to the reseller and reflected on his invoice.



Choose the customer and then click on the calendar to obtain statistics for desired period. The .csv file will have the following structure:

Account	From	To	Country	Description	Connect Time	Charged Time, min:sec	Charged Time, sec.	Charged Amount, USD
56.78.90.1	71886073902	380449313591	UKRAINE	Kiev Region	30-Apr-06 23:59:44	04:24	264	0.616
56.78.90.1	19190767456	420696017957	CZECH REPUBLIC	Proper	30-Apr-06 23:59:05	07:03	423	1.7625
56.78.90.3	60572324893	16049576339	CANADA	British Columbia	30-Apr-06 23:57:14	02:12	132	0.308
56.78.90.1	46890062001	380693412335	UKRAINE	Proper	30-Apr-06 23:56:50	03:52	232	0.58
200.45.23.1	46345311293	380442924858	UKRAINE	Kiev Region	30-Apr-06 23:56:26	02:48	168	0.392
56.78.90.1	82226061971	14257891107	UNITED STATES	Washington	30-Apr-06 23:55:52	02:32	152	0.076
56.78.90.1	573132489507	16047660320	CANADA	British Columbia	30-Apr-06 23:55:33	03:20	200	0.1
56.78.90.1	91438981472	420461329009	CZECH REPUBLIC	Proper	30-Apr-06 23:55:04	03:47	227	0.94584
200.45.23.1	31773453219	420971480263	CZECH REPUBLIC	Proper	30-Apr-06 23:54:41	03:11	191	0.79584
56.78.90.1	68027102122	380975904496	UKRAINE	Proper	30-Apr-06 23:53:42	08:49	529	1.3225
56.78.90.3	28478507435	420802725520	CZECH REPUBLIC	Proper	30-Apr-06 23:44:12	01:17	77	0.32084
56.78.90.1	15383396548	16042029917	CANADA	British Columbia	30-Apr-06 23:44:07	09:54	594	0.297

The following data is available:

- Account ID (if applicable; not available for a reseller’s xDRs, but available for xDRs of a reseller’s subcustomers).
- CLI or ANI (From).
- CLD or DNIS (To).
- Country and description of destination.
- Call start time.
- Charged time (in minutes:seconds).

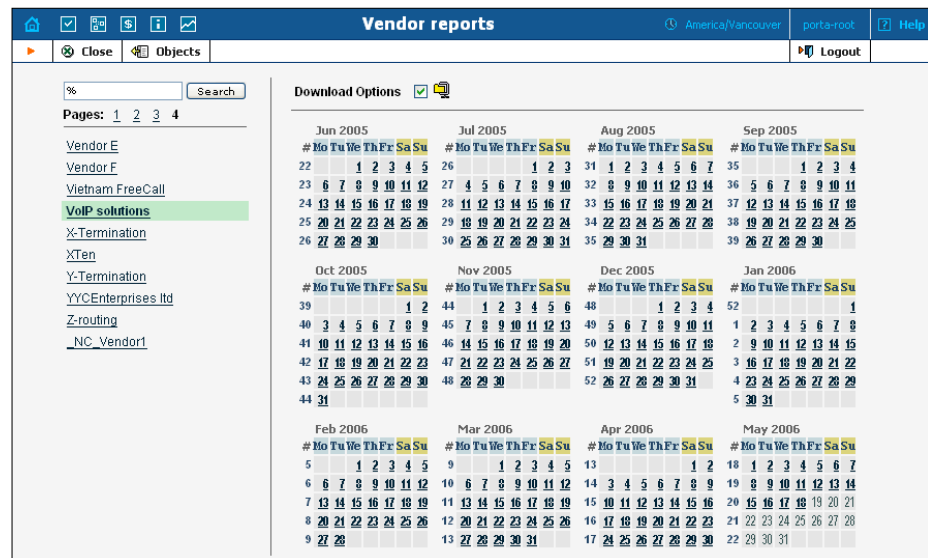
NOTE: When browsing data in Microsoft Excel, extra trailing zeros might be added, for example **345:37** (three hundred forty-five minutes and thirty-seven seconds) will be shown as **345:37:00**. Use proper call formatting in Excel to eliminate this problem.

- Charged time (in seconds).
- Charged amount.

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.

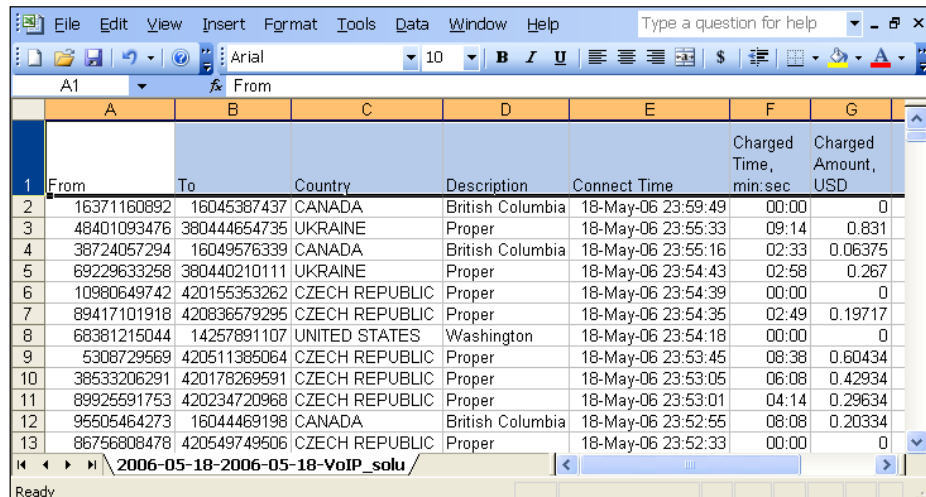
Vendor xDRs

It is very useful to have a detailed list of calls for a specific vendor, in case of disputes over the amount of terminated traffic. These statistics are calculated daily, and are split into files according to the vendor's billing period.



The screenshot shows the 'Vendor reports' interface. On the left, there is a search bar and a list of vendors: Vendor E, Vendor F, Vietnam FreeCall, **VoIP solutions** (highlighted), X-Termination, XTen, Y-Termination, YYCEnterprises ltd, Z-routing, and NC_Vendor1. The main area displays a grid of calendar views for each month from June 2005 to May 2006. Each calendar has columns for days of the week (Mo, Tu, We, Th, Fr, Sa, Su) and rows for days of the month. A 'Download Options' checkbox is checked. The interface also includes navigation buttons like 'Close', 'Objects', and 'Logout'.

Choose the vendor and then click on the calendar to download statistics for the desired period. The .csv file will have the following structure:



	A	B	C	D	E	F	G
1	From	To	Country	Description	Connect Time	Charged Time, min:sec	Charged Amount, USD
2	16371160892	16045387437	CANADA	British Columbia	18-May-06 23:59:49	00:00	0
3	48401093476	380444654735	UKRAINE	Proper	18-May-06 23:55:33	09:14	0.831
4	38724057294	16049576339	CANADA	British Columbia	18-May-06 23:55:16	02:33	0.06375
5	69229633258	380440210111	UKRAINE	Proper	18-May-06 23:54:43	02:58	0.267
6	10980649742	420155353262	CZECH REPUBLIC	Proper	18-May-06 23:54:39	00:00	0
7	89417101918	420836579295	CZECH REPUBLIC	Proper	18-May-06 23:54:35	02:49	0.19717
8	68381215044	14257891107	UNITED STATES	Washington	18-May-06 23:54:18	00:00	0
9	5308729569	420511385064	CZECH REPUBLIC	Proper	18-May-06 23:53:45	08:38	0.60434
10	38533206291	420178269591	CZECH REPUBLIC	Proper	18-May-06 23:53:05	06:08	0.42934
11	89925591753	420234720968	CZECH REPUBLIC	Proper	18-May-06 23:53:01	04:14	0.29634
12	95505464273	16044469198	CANADA	British Columbia	18-May-06 23:52:55	08:08	0.20334
13	86756808478	420549749506	CZECH REPUBLIC	Proper	18-May-06 23:52:33	00:00	0

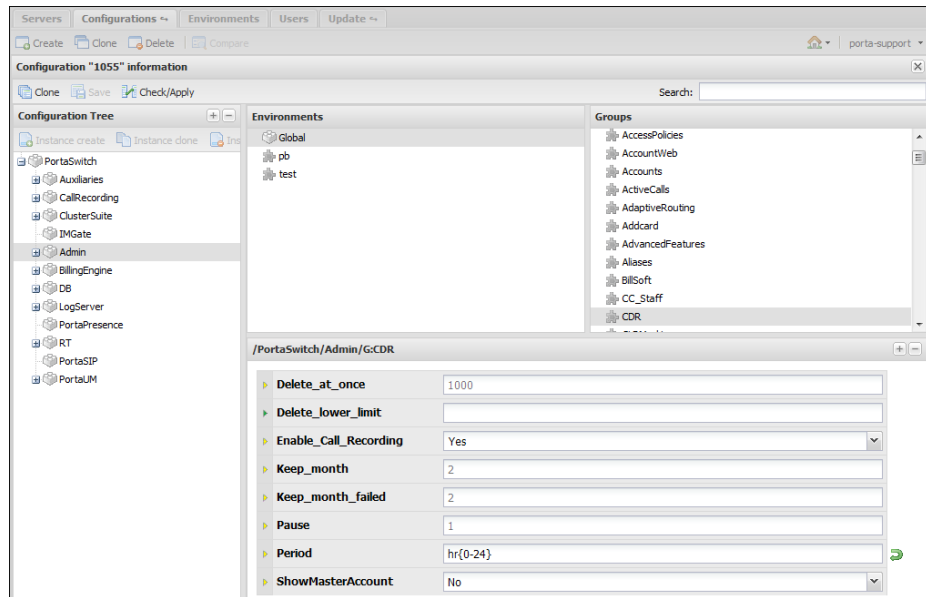
The following data is available:

- CLI or ANI (From)
- CLD or DNIS (To)
- Country and description of destination
- Call start time
- Charged time
- Charged amount

xDR Cleanup Procedure

xDRs (do not mix with billing statistics) are stored in the Main database and replicated to the Replica database. In order to prevent system overgrowing, xDRs are periodically cleaned up by performing a special task. To correctly reflect a charge history a special “Cleanup” xDR exists that accumulates the charged amounts and times from all the xDRs that have been cleaned up for every 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 the xDRs that are more than 2 months old, though these options are configurable to meet special requirements via the Configurator.



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.

Cost / Revenue Statistics

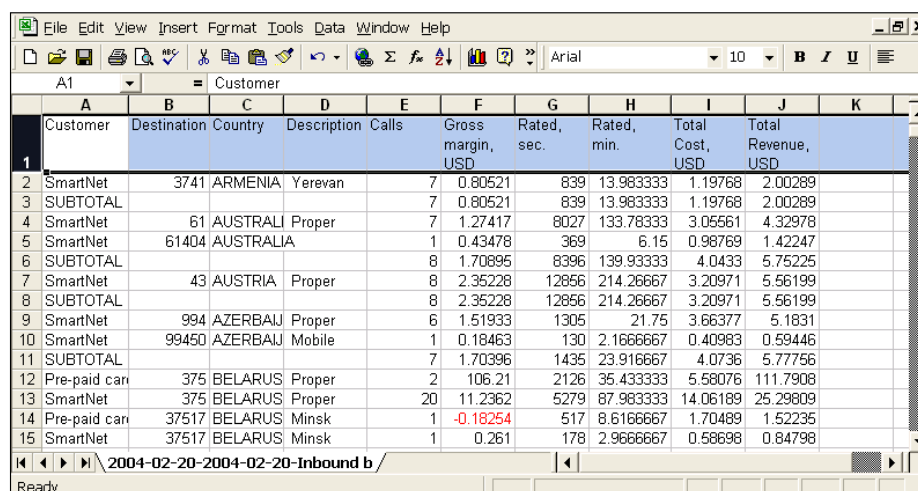
These are essential tools for monitoring the growth of your business. Working with different partners, different currencies and different prices, it is very easy to make mistakes and carry out non-profitable calls.

Cost/Revenue reports allow you to monitor this and stay out of trouble. To access Cost/Revenue reports, follow the **Cost/Revenue reports** link from the **Statistics** section of the main menu.



There are five variants of pre-calculated reports:

- By customer and destination, subtotal per country.
- By customer and destination, subtotal per customer.
- By vendor and destination, subtotal per country.
- By vendor and destination, subtotal per vendor.
- By destination.



	A	B	C	D	E	F	G	H	I	J	K
	Customer	Destination	Country	Description	Calls	Gross margin, USD	Rated, sec.	Rated, min.	Total Cost, USD	Total Revenue, USD	
2	SmartNet	3741	ARMENIA	Yerevan	7	0.80521	839	13.983333	1.19768	2.00289	
3	SUBTOTAL										
4	SmartNet	61	AUSTRALI	Proper	7	1.27417	8027	133.783333	3.05561	4.32978	
5	SmartNet	61404	AUSTRALIA		1	0.43478	369	6.15	0.98769	1.42247	
6	SUBTOTAL										
7	SmartNet	43	AUSTRIA	Proper	8	2.35228	12856	214.26667	3.20971	5.56199	
8	SUBTOTAL										
9	SmartNet	994	AZERBAIJ	Proper	6	1.51933	1305	21.75	3.66377	5.1831	
10	SmartNet	99450	AZERBAIJ	Mobile	1	0.18463	130	2.1666667	0.40983	0.59446	
11	SUBTOTAL										
12	Pre-paid cari	375	BELARUS	Proper	2	106.21	2126	35.433333	5.58076	111.7908	
13	SmartNet	375	BELARUS	Proper	20	11.2362	5279	87.983333	14.06189	25.29809	
14	Pre-paid cari	37517	BELARUS	Minsk	1	-0.18254	517	8.6166667	1.70489	1.52235	
15	SmartNet	37517	BELARUS	Minsk	1	0.261	178	2.9666667	0.58698	0.84798	

The following data is available:

- Name of customer/vendor.
- Destination prefix.
- Country and description of destination.
- Total number of successful calls.
- Gross margin – The difference between total revenue and the cost for this destination.
- Rated, sec. – The sum of the time charged for all calls to this destination in seconds.
- Rated, min. – The sum of the time charged for all calls to this destination in minutes. This is the same value as in the previous column, only expressed in different units. Thus, for instance, if *Rated, sec* contains **180**, here **3.0** will be shown.
- Total Cost – Summary of the amount charged for all the vendor's CDRs for this destination. If the vendor uses a currency different than your base one, this will be converted using the current exchange rate.
- Total Revenue – Summary of the amount charged for all the customer's CDRs for this destination. If the customer uses a currency different than your base one, this will be converted using the current exchange rate. Note that, in the case of resellers, we use the value of the reseller's CDR, not the subcustomers, since your revenue is what gets invoiced to the reseller.

In addition, you can use the **Custom Query** report, which is identical to the **Custom Query** report in the ASR statistics.

Encrypted Passwords

Security matters for every ITSP company since it is important to supply

customers, especially corporate ones, with quality and with protected services. For this purpose, there is an encrypted password functionality that increases the level of protection for passwords stored in a database. This protection measure will protect users' passwords if or when an unauthorized person obtains access to the database.

Passwords are encrypted using salt (random data used to randomize output after encryption). Salt complicates the process of deciphering the encrypted password since each salt value is unique, even for two identical passwords.

Passwords encrypted with salt are stored in a database, and keys that are used for encrypting or decrypting the passwords are stored in a file on servers with Billing Engine and PortaSIP® instances. So encrypted passwords and keys are stored in different places and for this reason, it is difficult for passwords to be cracked, even if an attacker obtains access to the database.

To activate the encrypted password functionality the administrator should go to the configuration server web interface and set the **Security.EncryptPasswords** option to **Yes** on the configuration server web interface. After saving this configuration, the system will encrypt passwords for new accounts and passwords that will be changed for existing accounts.

Specifics of this functionality:

- An administrator can initiate the key change procedure during which an old encryption key 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 services as usual since there is no downtime.
- If an administrator schedules backup, the keys won't be included in this backup.

The encrypted password functionality increases the security level for users' passwords stored in a database.

8 ■ How to ...

Log in to PortaBilling® After Installation

There is one default environment pre-created, which contains a single user “pb-root” with the password “pb-root” as well. After you login for the first time, you will be requested to change the password – please remember the new password you choose.

Terminate a Customer, Including All His Accounts

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, use the **Terminate** button on the **Customer Info** toolbar. 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 such a customer is by using **Advanced Search** with the “closed status” filter.

Reversible 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 him / her. 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.

Locate the h323-conf-id for a Call

Unique call IDs are extremely important for troubleshooting. They allow you locate a call in the database and find call information in the billing engine logs or RADIUS detail files. Finally (and this is the most important

thing), when reporting call problems to the PortaOne Support team or your business partner, a call ID helps to prevent confusion and allows the problem to be solved quickly.

In order to find the h323-conf-id for a call, do the following:

1. From the main menu, go to **Trace Session**.
2. Using the CLD and from/to date filters, make sure the required call is shown in the list of calls on the page.
3. In the row containing call information, click the  **View** icon (far left column). You will be shown detailed information about the call.
4. The first field on the page should be the h323-conf-id, which is a combination of four hexadecimal numbers separated by spaces, e.g. 5FF7F6D1 715E02C6 A40990F3 C823E27E.

Troubleshoot an Incorrectly Billed Call

1. Make sure that someone in your organization is subscribed to the PortaBilling® mail alerts (especially “Missing critical information”). This will help you to detect problems early.
2. Find the h323-conf-id for this call. This is a unique ID (a string of four hex numbers) generated by the gateway when the call was started, which will help you to exactly identify this particular call among all the others. You can find this ID by doing one of the following:
 - Looking at the statistics on your gateway.
 - In the **BE Log Viewer** section of the web interface you can find a list of recent calls (a select menu allows you to see all calls made in the past 5, 10, 30, etc. minutes).
3. Use **BE Log Viewer** on the web interface to find all of the information about call processing there.

NOTE: Normally we receive information about each of the call legs separately, so it is necessary that you check the log entry regarding the processing of **all** call legs and the final call clean-up.

4. The following are typical error situations:
 - **Call was not billed at all.** It was considered an “unresolved” call, because we did not detect that it went to any of the vendors. Check that you have correctly defined connections to your vendors.
 - **Call was billed only to vendor, not to an account or customer** (and you received an email alert). The billing engine needs an Account ID in the User-Name attribute to correctly identify the account and customer. Check the

logs to see what was in the User-Name attribute. Typical situations are:

- **Incorrect value in User-Name** (for example, phone number instead of IP address of the remote gateway). Cause: the required application was not used to handle the call. Remedy: check that the application is configured and associated with the corresponding *incoming* dial-peer.
 - **On some of the call legs there is a correct value in the User-Name, while on others there is an IP address.** When a call is originated on gateway A and terminated on gateway B, then a “real” username appears only in the accounting from gateway A. In gateway’s B accounting the username will be the IP address of gateway A, because gateway A had to authenticate itself before being able to make the call. This is a perfectly normal situation. PortaBilling® recognizes this, and will replace a username identical to the node’s ID (IP address) with a “real” username from the other call legs.
 - **Value in User-Name is correct, but reports “Did not find account/customer”.** Check the rating table for the account’s product. It may be that, even if you have such an account, the rating table in the account’s product is incorrect (it does not include the service consumption point where you are trying to use the service).
- **Call was billed, but the phone number in the billing is incorrect (not in E.164 format) and there is a “Mismatch in rates or destinations” error.** Cause: missing or incorrect translation rule on connection to the vendor. Remedy: assign a translation rule; read the **Handle technical prefixes and numbering formats** section in the **Unified PortaSwitch Handbook Collection**.
 - **Call was billed and the phone number in the billing is correct, but there is a “Mismatch in rates or destinations” error.** Cause: missing rate for this phone prefix in the tariff. Remedy: create a rate for this destination in the tariff.
5. If you are unable to solve the problem by yourself, submit a problem report to support@portaone.com. Please make sure that you include the following in your email:
- A detailed description of the call flow and what seems to be incorrect.

- H323-conf-id of the call.
- Relevant items from the billing engine log.

Make the Periodic Payments Tab Appear in the Customer / Account Info

The Periodical payments tab will only be shown if the rest of the system is configured to accept payments. This requires:

- At least one payment system defined in the **Payments** section; this payment system must have the **Recurring** option enabled.
- A payment system associated with the currency of this account or customer.
- For accounts, the **E-commerce enabled** flag must also be switched on.
- On the **Payment info** tab, the preferred payment method must be set and information about the credit card entered.


For subcustomers or accounts belonging to subcustomers, the requirements are the same, except that the payment systems are defined in the reseller information, not in Payments section.

Make the Custom Fields Tab Appear in the Customer / Account Info

Custom fields should be initially set up using the **Custom Fields** tab on the **Web Interface** page. Only when configured the corresponding tab will appear in the customer / account info. Using custom fields, you can store a set of extra attributes (e.g. driver's license ID or tax code) that supplements the standard PortaBilling information.

For each new custom field, the following attributes must be set:

Field	Description
Object	Defines whether the custom field applies to the Customer or the Account.
Name	The descriptive name of the field. This is the name that will be displayed next to the custom field on the Customer/Account Info page.

Type	<p>Defines the type of field:</p> <ul style="list-style-type: none"> • Text – Basic single-line input field. • Number – Input field used to store and validate numerical values. • Date – Field type used to store dates. • Date & Time – Custom field that stores dates with a time component. • List – Single select list with a configurable set of options.
Properties	<p>Enables you to customize properties of the field that define its form, appearance, or value. These properties are specific to the field type.</p> <p>Click Properties or the  Wizard icon to invoke the wizard. This will enable you to define a new field format or change an existing one and to specify the default value a custom field should have.</p>
Default	<p>Read-only attribute which must be specified in the Properties attribute.</p>
Mandatory	<p>Defines the mandatory status of the field.</p>



You can delete a custom field at any time. All records of its values will then also be deleted.

Make a Custom Report from PortaBilling®

PortaBilling® provides you with an open data model. An ER-diagram of the database structure is not included in this document, but may be received from the PortaOne support team upon request. If you want to prepare custom reports on your workstation or a non-PortaBilling® server, you will need to do the following:

- Make sure the remote computer has database drivers installed to access the PortaBilling® database. Normally you would use native MySQL connectivity on Unix-based hosts and ODBC on Windows-based hosts.
- For any data-mining solutions (extracting data from the database), use **only** the replica database.
- Use a tool like Crystal Reports, Microsoft Access or some custom application to retrieve data from the database, process it and submit it to the user.

Use ODBC to Connect to PortaBilling®

ODBC (Open Database Connectivity) provides a way for client programs to access a wide range of databases or data sources. If you need extended customized reporting not available in PortaBilling®, you can do this using external tools such as MS Access or Crystal Reports.

Create a MySQL user to be used for reports

1. Log in into your PortaBilling® replica DB server using ssh.
2. Start the MySQL command tool.

```
andrew@demo:/home/porta-admin$mysql -u root mysql
Reading table information for completion of table and column
names
You can turn off this feature to get a quicker startup with -A
Welcome to the MySQL monitor. Commands end with ; or \g.
Your MySQL connection id is 42122 to server version: 4.0.17-log
Type 'help;' or '\h' for help. Type '\c' to clear the
buffer.
mysql>
```

3. Create a new user using the GRANT command.

```
mysql> grant ALL PRIVILEGES on `porta-billing`.* to
`reports`@'192.168.0.5' identified by `pod23uk`;
Query OK, 0 rows affected (0.02 sec)
```




NOTE: The command above will permit access to all of the tables in the database. It is provided only as an example. You can modify it according to your actual needs, e.g. give **read-only** access to **some** tables only.




4. Flush the privileges.


```
mysql> FLUSH PRIVILEGES;
Query OK, 0 rows affected (0.07 sec)
```

Installing the MySQL ODBC driver

1. Download and run the installation package from:
<http://www.mysql.com>
2. Click **Next** on the information pages.







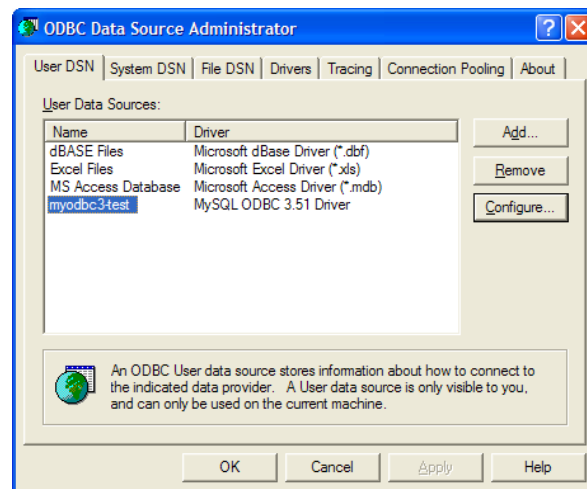
3. Click **Finish**.

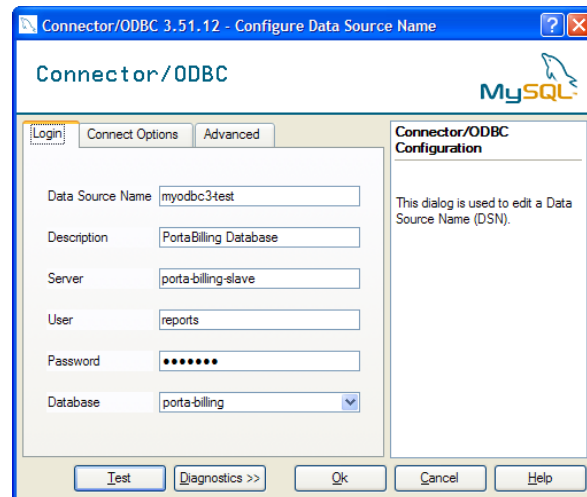
Configuring ODBC

Before configuring the data source, create a MySQL user on replica DB server with read-only permissions. Please examine the following document on how to add new user accounts to MySQL:

http://www.mysql.com/doc/en/Adding_users.html

1. Control panel → Administrative tools → Data sources (ODBC).

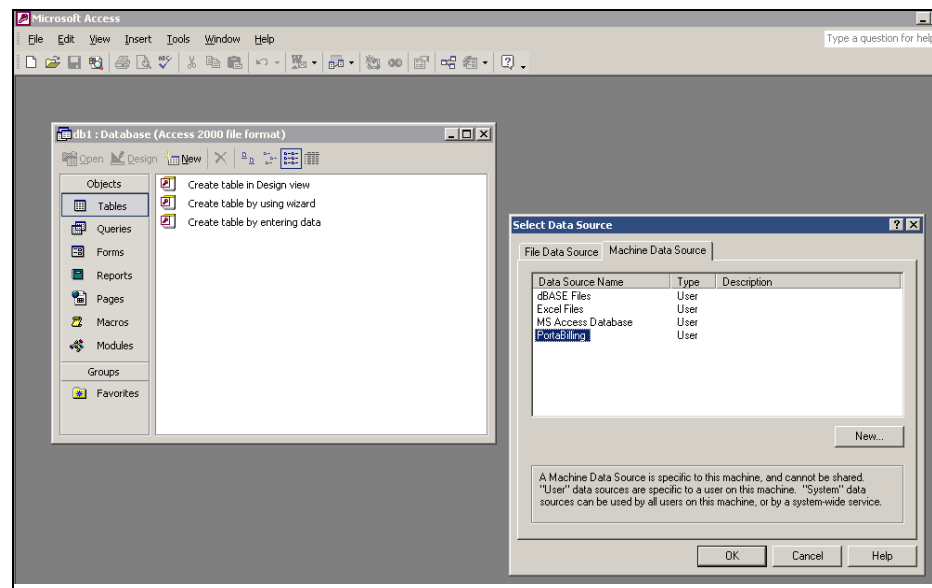




2. Select myodbc3-test and click **Configure**. Fill in the configuration form. Important parameters include:
 - **Host/Server Name (or IP)** – Hostname (or IP address) of your **replica DB** server.
 - **Database Name** – **porta-billing**.
 - **User** – Username of the MySQL user you have created for reporting purposes.
 - **Password** – Password of the MySQL user you have created for reporting purposes.
 - **Port** – The port on which the database service is accessible; enter **3307** here.

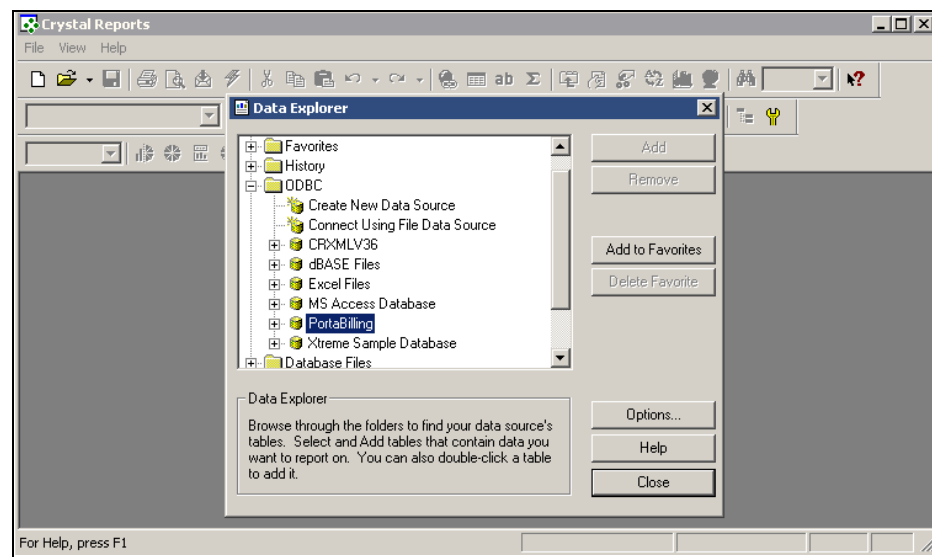
NOTE: This port number differs from the one used by default.

Using ODBC



In MS Access

1. Create a blank database.
2. Right-click in the table design view, and select **Link tables**.
3. Choose **ODBC databases** from the **Files of type** list.
4. Select **Machine data source**.
5. Select **PortaBilling** and click **OK**.
6. Select the desired tables.



In Crystal Reports

1. Create a new report.

2. In **Data Explorer**, open **ODBC** branch.
3. Select **PortaBilling**.
4. Select the desired tables.

Force PortaBilling® to Disconnect after a Customer Calls over his Credit Limit

There is no need for PortaBilling® to do this, as the gateway is able to by itself. When the gateway authorizes an account to make a call in PortaBilling®, PortaBilling® returns a maximum credit time in the case of a successful authorization. When the gateway connects the call, it starts a timer; when the timer hits zero, it automatically disconnects the call. PortaBilling® sends **two** RADIUS attributes related to the maximum call duration. The `h323-credit-time` attribute is the standard one (announced credit time); there is also a custom `h323-ivr-in = DURATION` attribute (actual credit time). The value for the second attribute is computed with billing tricks taken into account. Thus, in the case of \$10 available funds, a \$0.10/min rate and a 10% post-call surcharge, this attribute will be 5400 seconds (90 minutes), while the `h323-credit-time` will be 6000 seconds (100 minutes).

If the application supports only the `h323-credit-time` attribute (e.g. Cisco default debit card application), then this time will be announced to the end user and used to disconnect the call, and `h323-ivr-in = DURATION` will be ignored.

Create Accounts to be Used for SIP Services

There are no special requirements as to how such accounts should be created. You use the same interface to create and manage accounts for all services supported by PortaBilling® (H323, SIP). Thereafter accounts can use H323, SIP or SIP & H323 services, depending on their product's rating table. So if you plan for accounts with a certain product to be able to login to the SIP server and make outgoing calls, be sure to include the PortaSIP® node with the appropriate tariff in the rating table (**Usage Charges** tab) for this product.

Integrate PB Logins to Your Website

You can include the login form on your website using the following HTML code:

```
<form name="log" method="post"
action="https://pb.mycompany.com/index.html">
  Login <input type="text" name="pb_auth_user" value="">
```

```
    Password <input type="password" name="pb_auth_password"
value="">
    <input type="submit" value="Login">
    <input type="hidden" name="redirectOnLoginError"
value="http://www.myweb.com">
</form>
```

“redirectOnLoginError” - in case of unsuccessful logging in, PortaBilling® will redirect the browser to the URL specified in the value attribute.

Add Your Company Name and Logo to Every PortaBilling® Web Page

Using the brandpane feature you can customize your system on a per-environment basis. 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.tmpl,’ 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=/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=/home/demoroot/tmp/bp
```

2. Edit files, add / remove images, styles, etc.:

```
nano /home/demoroot/tmp/bp/default/admin/branpane.tmpl
```

3. Import brandpane templates back into the database:

```
impexp_templates.pl -i -brandpane=/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:

```
impexp_templates.pl -d -t=122
```

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>
        <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.

Add 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

% }

```

Automatically Purge Old xDRs to Reduce the Size of the Database

Database size continually increases, thus affecting system performance: longer PDD (Post Dial Delay), longer response time on web interface pages related to xDRs (rate upload, xDR browser, invoicing) etc. Large xDRs tables also increase database query execution time. xDR Cleanup procedure is periodically executed on your system to ensure the smooth operation. By default the execution is scheduled for every day during your off-peak hours.

The maximum period for keeping xDRs depends on your requirements. We recommend keeping xDRs for the period that corresponds to the maximum period you need to generate reports plus an additional month. Thus, if you only need monthly reports, keeping the xDRs for 2 months will be sufficient.

9 ■ 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. 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 C. 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 D. 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.

ACD – Automated Call Distribution

AD – Abbreviated Dialing

ADSL – Asymmetric Digital Subscriber Line

AIS – Automated Information System

ANI – Automatic Number Identification

API – Application Programming Interface

ART – Audible Ringing Tone

AS – Answer Supervision

ASP – Application Service Provider

ASR – Answer Seizure Ratio

ATA – Analog Telephone Adapter

AVVID – Architecture for Voice, Video and Integrated Data

B

B2B – Business to Business

B2C – Business to Consumer

BH – Busy Hour

BHCA – Busy Hour Call Attempts

BHCC – Busy Hour Call Completion

BIOS – Basic Input/Output System

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.

CALEA – Communications Assistance for Law Enforcement Act

CDR – Call Detail Record

CGI – Common Gateway Interface

CHAP – Challenge Handshake Authentication Protocol

CLE – Customer Located Equipment

CLEC – Competitive Local Exchange Carrier

CLI – Command Line Interface

CLI, CLID – Calling Line Identification

CMS – Call Management Services

CPE – Customer Premise Equipment

CRM – Customer Relationship Management

D

DCE – Data Communications Equipment

DHCP – Dynamic Host Configuration Protocol

DID – Direct Inward Dialing

DMZ – Demilitarized Zone

DNS – Domain Name System

DSL – Digital Subscriber Line

DTE – Data Terminal Equipment

DTMF – Dual-tone multi-frequency. The use of two simultaneous voice-band tones for dialing (such as touch-tone).

E

E-1 – Europa-1, the European equivalent to T1 (but higher bandwidth).

E911 – Enhanced 911

ENUM – Electronic Numbering

F

FAQ – Frequently Asked Questions

FCC – Federal Communications Commission

FoIP – Fax over IP

FQDN – Fully qualified domain name. The full name of a system, rather than just its hostname. For example, test is a hostname, while test.portaone.com is an FQDN.

FXO – Foreign Exchange Office

FXS – Foreign Exchange Station

G

GK – Gatekeeper. The entity that maintains a registry of devices in an h323 network.

GMT – Greenwich Mean Time

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

HDD – Hard Disk Drive

HTML – Hypertext Markup Language

HTTP – Hypertext Transfer Protocol

I

ICMP – Internet Control Message Protocol

IM – Instant Messaging

IMAP – Internet Message Access Protocol

IMS – IP Multimedia Subsystem

IOS – Internetworking Operating System

IP PBX – Internet Protocol Private Branch Exchange

ISDN – Integrated Services Digital Network

ISP – Internet Service Provider

ISV – Independent Software Vendor

ITSP – Internet Telephony Service Provider

IVM – Instant Voice Messaging

IVR – Interactive Voice Response

IXC – Inter–Exchange Carrier

J

JMS – Java Messaging Service

JVM – Java Virtual Machine

L

LAN – Local Area Network

LBS – Location Based Services

LCR – Least Cost Routing

LDAP – Lightweight Directory Access Protocol

LEC – Local Exchange Carrier

Long pound calls – see *Multi-session call*

M

MAC – Media Access Control

MAC address – The physical address of a device on Ethernet.
MG – Media Gateway
MGCP – Media Gateway Control Protocol
MIME – Multi-Purpose Internet Mail Extension
MMS – Multimedia Messaging Service
MOD – Movies on Demand
Multi-session call – A call in which several outgoing calls to different destinations can be made under the same incoming call to the IVR.

N

NANP – North American Numbering Plan
NAS – Network Access Server
NAT – Network Address Translation
NBH – Network Busy Hour
NGN – Next Generation Network
NIC – Network Interface Card
NNX – The current general configuration for exchange codes within each area code.
NOC – Network Operations Center
NPA – Numbering Plan Area
NTP – Network Time Protocol

O

OEM – Original Equipment Manufacturer
OLTP – Online Transaction Processing
OS – Operating System
OSD – On-Screen Display

P

PABX – Private Automatic Branch Exchange
PAP – Password Authentication Protocol
PBX – Private Branch Exchange
PDF – Portable Document Format
PIN – Personal Identification Number
PnP – Plug and Play
POP – Point of Presence
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
PSAP – Public Service Answering Point
PSTN – Public Switched Telephone Network

Q

QoS – Quality of Service

R

RADIUS – Remote Authentication Dial-In User Service
RAM – Random Access Memory
RAS – Registration, Admission and Status protocol
RDBMS – Relational Data Base Management System
RFP – Request for Proposal
RFQ – Request for Quote
RJ – Registered Jack
RS-232 – Recommended Standard 232 (physical interface between DCE and DTE)
RTP – Real-Time Transfer Protocol
RTCP – Real-Time Control Protocol
RTSP – Real-Time Streaming Protocol

S

SBC – Session Border Controller
SCSI – Small Computer Systems Interface
SDK – Software Developers Kit
SDP – Session Definition/Description Protocol
SG – Signaling Gateway
SIGTRAN – Signaling Transport
SIM – Subscriber Identify Module
SIM Card – Subscriber Identify Module Card
SIP – Session Initiation Protocol
SIP-T – Session Initiation Protocol for Telephony
SLA – Service Level Agreement
SMP – Symmetric Multiprocessing
SMS – Short Messaging Service
SMTP – Simple Mail Transfer Protocol
SOAP – Simple Object Access Protocol
SOHO – Small Office/Home Office
SQL – Structured Query Language
SS7 – Signal System 7

T

T.37 – Standard for store-and-forward FoIP (Fax over IP).
T.38 – Standard for real-time FoIP (Fax over IP).
TAPI – Telephony Application Program Interface
TCL – Tool Command Language. A scripting language used for gateway products both internal and external to the Cisco IOS software code.
TCP – Transmission Control Protocol
TCP/IP – Transmission Control Protocol/Internet Protocol
Telco – A telephone company, such as LEC.
TFTP – Trivial File Transfer Protocol
TTS – Text-to-Speech
TVoIP – TV over IP

U

UA – User Agent
URI – Uniform Resource Identifier
URL – Uniform Resource Locator
USB – Universal Serial Bus

V

VAD – Voice Activity Detection
VAR – Value-Added Reseller
VLAN – Virtual LAN
VOD – Video on Demand
VoiceXML – Voice Extensible Markup Language
VoIP – Voice over Internet Protocol

W

WAN – Wide Area Network
WiFi – Wireless Fidelity
WPA – WiFi Protected Access
WWW – World Wide Web

X,Y,Z

xDR – Extensible Detail Record
XML – Extensible Markup Language

APPENDIX F. Hotlining Parameters

Parameter	Description
fui=url	This enables the hotlining feature.
hotline_url= <i>http://thewebsite.com</i>	This points to the redirection website.
hotline_denied=1	This enables hotlining even for blocked, suspended or expired accounts