# Cable Modem Load Balancing for Cisco Series uBRs

**Table of Contents**



## Load Balancing Feature Overview

Load Balancing on the Cisco CMTS allows service providers to optimally use both downstream and upstream bandwidth, enabling the deployment of new, high-speed services such as voice and video services. This feature can also help reduce network congestion without requiring expensive node-splitting.  Congestion is due to the uneven distribution of cable modems across the cable network and due to different usage patterns of individual customers.

By default, the Cisco CMTS platforms use a form of load balancing that attempts to equally distribute the cable modems to different upstreams when the cable modems register. You can refine this form of load balancing by imposing a limit on the number of cable modems that can register on any particular upstream, using the cable upstream admission-control command.

However, this default form of load balancing affects the cable modems only when they initially register with the Cisco CMTS. It does not dynamically rebalance the cable modems at later times, such as when they might change upstream channels in response to RF noise problems, or when bandwidth conditions change rapidly because of real-time traffic such as Voice over IP (VoIP) and video services. It also does not affect how the cable modems are distributed among downstream channels. This form of load balancing allows user-configurable parameters to provide optimization on a per CMTS basis. This eliminates the use of a provisioning system for modem distribution.

## Load Balancing Options

There are 2 forms of load balancing; static and dynamic.

## Static Load Balancing

Static load balancing is registration-based where CMs only balance when registering with the CMTS. There is also a passive mode for troublesome CMs**.**

For modems that do not respond to normal static load balancing, you can configure a passive type of static load balancing on a per-OUI basis. The CMTS ignores the RNG-REQ from modems attempting to register on an overloaded DS or US. This forces the modem to search for a channel that it can register on.

```
cable load-balance group <g> protocol upstream {direct | override | passive}
```

This is a hidden command and "override" is the default setting.

Below are listed three subsets of static load balancing:

**direct** - just tell the modem it is on the correct US channel. This should work if modems follow the DOCSIS protocol as written, but unfortunately many don't.

**override** - Use "channel override TLV fields in the ranging response messages to tell modems which US (or DS) they are supposed to use. This is the default setting for static load balancing. Unfortunately, some modems don't support this either.

**passive** - ignore the modem until it shows up where we want it. This can cause problems because modems will keep trying and increasing power levels, which can ultimately cause laser clipping if many modems do it at the same time (e.g., after the CMTS is rebooted).

## Dynamic Load Balancing

Dynamic load balancing moves CMs “on the fly” while online. Modems don’t need to go offline (re-register) to be balanced between US ports, but downstream channel changes force modems to reacquire DS sync, range, and re-register.

Upstream load balancing does NOT require modems to go offline to change upstream ports. This is vendor specific with different types of modems. With 1.0 and some 1.1 modems, they may re-register when executing an upstream channel change.

## Load Balancing Methods

There are three methods of load balancing, which are modem, service flow, and utilization.

## Modem Method

The modem method of load-balancing uses the number of active cable modems on an interface to determine the current load. This is a form of distribution-based load balancing, in which the absolute numbers of modems are used to determine whether interfaces are load balanced.

This method does not take into account the amount of traffic flowing through the cable modems, but the system does take into account the relative bandwidth of the channels being used, so that channels with higher bandwidths are allocated higher numbers of cable modems. This means that when interfaces are using different channel widths or modulation profiles, the system can assign different numbers of cable modems to the interfaces to achieve a balanced load.

For example:
Channel widths— If two upstreams are being load balanced with one US configured with a channel width of 1.6 MHz and the other US is configured for 3.2 MHz, the CMTS allocates twice as many CMs to the second US, because its channel width is twice as large as the first.

Modulation profiles— If one downstream is configured for 64-QAM and the other downstream is configured for 256-QAM, the Cisco CMTS allocates a proportionately larger number of cable modems to the second downstream so as to achieve a balanced load.

When both the channel width and different modulation profile are set differently on two interfaces, the system calculates a "weight" value to use as a guide to determine the relative bandwidths of the interfaces.

## Service Flow Method

This method uses the number of active service flow IDs (SFIDs) on an interface to determine the current load. This is a form of distribution-based load balancing, in which the absolute number of service flows are used to determine whether interfaces are load balanced. In some cases, a CM could have many simultaneous DOCSIS 1.1 service flows.

This method does not take into account the amount of traffic flowing on each SFID, but the system does take into account the relative bandwidth of the channels being used, so that channels with higher bandwidths are allocated higher numbers of SFIDs. This means that when interfaces are using different channel widths or modulation profiles, the system can assign different numbers of SFIDs to the interfaces to achieve a balanced load.

When both the channel width and different modulation profile are set differently on two interfaces, the system calculates a "weight" value to use as a guide to determine the relative bandwidths of the interfaces.

## Utilization Method

The utilization method uses an interface's current percentage of utilization to determine the current load. This method uses the amount of traffic being sent over an interface, in the form of the percentage of total bandwidth being used. The system takes into account the relative throughput and bandwidth (as determined by the modulation profiles and channel widths) of each interface when evaluating the load on those interfaces.

For example, if two upstreams are being load-balanced using the utilization method, and the first upstream has twice the bandwidth of the second upstream, the two upstreams are considered balanced when they reach the same percentage of utilization. The first upstream is carrying more traffic than the second upstream because it has a larger capacity for traffic, but the percentage of utilization will be relatively the same.

## Load-Balance Groups (LBG)

Load balance groups have a few parameters that can be configured.

- **cable load-balance group n method {modem, service-flows, utilization}**
- **cable load-balance group n interval {seconds}**
- **cable load-balance group n threshold load minimum {min-value}**
- **cable load-balance group n threshold load {load-value} | enforce {enforce-value}**
- **cable load-balance group n policy ugs**
- **cable load-balance group n threshold ugs {ugs-value}**
- **cable load-balance group n threshold stability {percentage}**

✦ cable load-balance group n method {modem, service-flows, utilization}

**Method –**

These three methods were discussed earlier.

✦ cable load-balance group n interval {seconds}

**Interval –**

Valid range of 1-1000 with a default of 10 seconds. This is the interval of time that the load-balance algorithm will check to see if the user-configurable thresholds have been exceeded. Once the algorithm is activated, it will continue until the load is exactly balanced or check to see if the load is within the thresholds again at the next interval.

✦ cable load-balance group n threshold load minimum {min-value}

**Threshold Load Minimum –**

Valid range is 1 to 100 with a default of 5. This specifies the difference in number of modems between any two of the ports assigned to that LBG. This number is checked each interval period. If the difference between any two ports is less than the minimum value configured, no load-balancing will be performed. If equal to or greater than the configured value, **AND** the thresholds set in load (for static load-balancing) or enforce (for dynamic load-balancing) are not met, load-balancing will be performed. **Note**: This command is not valid for utilization method.

✦ cable load-balance group n threshold load {load-value}

**Threshold Load –**

Value in percentage (1% – 100%) with a default of 10%. This specifies the maximum load difference that can exist between any two ports of a LBG. Load value is used to perform load-balancing statically as modems register. If the minimum threshold is exceeded, the load threshold is checked to determine if load-balancing should start.

✦ cable load-balance group n threshold load {load-value} | enforce {enforce-value}

**Threshold Load x Enforce –**

Valid range of (<current load value> to 100%) with the default being equal to the <current load value>. This enables dynamic (post-registration) load-balancing and specifies the maximum load difference that can exist between any two ports of a LBG. If the minimum threshold is exceeded, the enforce threshold is checked to determine if load-balancing should start. **Note**: *This command is optional and used to enable dynamic load balancing. If not specified only static load balancing will occur.*

✦ cable load-balance group n policy ugs

**Policy ugs -**

The default is to enable modems with active ugs flows to participate in load balancing after a certain threshold is reached, which is 70% by default as listed below. Care needs to be taken if one chooses not to balance modems with ACTIVE UGS flows [phone calls]. It's possible to disable this feature by using the NO form of this command.

✦ cable load-balance group n threshold ugs {ugs-value}

**Threshold ugs -**

Default of 70%. When the modem count is unbalanced, modems will move, but not UGS CMs. If the reserved bandwidth is above 70% while the modem count is unbalanced, random modems will move whether they have UGS flows or not until the reserved bandwidth is below 70% and the modem count is balanced within its configured thresholds. Reserved bandwidth is UGS, RTPS, and CMs with guaranteed US rates.

✦ cable load-balance group n threshold stability {percentage}

**Threshold Stability -**

Default = 50% of good station maintenance (SM) messages. If the % of good SM gets below this setting, then the CMTS will mark the US unstable. "unstable" means that too many ranging requests got lost. The configured % is the "acceptable" rate at which an interface is considered stable in a 1 min period. LB will not touch unstable ports. It may move CMs away from an unstable port, but it will not move CMs to it. It will not move all CMs away from an unstable interface - at best it will try to achieve balanced load by moving CMs away from it. It's the responsibility of applications such as spectrum management or redundancy to control parameters if an interface performs badly, and if necessary to shut it down. Load balancing does not attempt to "shadow" those other applications, and it would be a bad idea to try. Such an attempt might cause system instabilities by interfering with those other applications. You can see current percentage values with the CMTS command, "**show cable load-balance internal**".

To enable the Load Balancing feature on a Cisco CMTS, you first must:

1. Create and configure a load-balance group, which specifies how load balancing should be performed.
2. Assign cable interfaces to the load-balance group, at which point the Cisco CMTS begins performing load balancing on those cable interfaces.

You can use separate load-balance groups for upstreams or downstreams, or you can use the same load-balance group for both upstreams and downstreams. However, all cable interfaces in a load-balance group must share the same physical RF connectivity.

Also, the same load-balance group must be used for all downstreams or upstreams that share RF connectivity and that are participating in load balancing. You cannot distribute downstreams or upstreams that share physical connectivity across multiple load-balance groups.

If you assign downstreams and upstreams to different load-balance groups, the Cisco CMTS performs load balancing independently on the upstreams and downstreams.

If both DSs & USs are assigned to the same LBG, the CMTS attempts to balance both DS and US load with DS taking precedence. However, if the US load on the target interface is higher than the US load on the source interface, no modem will be moved. Or, in other words, LB will only move a modem from a higher utilized DS channel (interface) to another if it finds an US channel with lower load than the originating US channel. This is true if the US channels are configured into LB groups (which is recommended). If only DS channels are configured, the system will move modems from the interface with the lower load to the interface with higher load and will ignore US channel load.

## Load Balancing (LB) Restrictions

The load-balancing algorithms assume a relatively even distribution of usage among modems. Load balancing is done on cable modems in real-time, using current load-usage statistics. Load balancing can not be performed according to the time-of-day or using a schedule.

Particular cable modems cannot be selected to be moved for load balancing, although cable modems can be excluded from load balancing operations altogether on the basis of their MAC address or organization unique identifier (OUI).

Load balancing can be done only on a per chassis basis—all interfaces in a load-balance group must reside in the same chassis.

A downstream or upstream can belong to only one load-balance group  You can configure only one load-balance group per shared physical domain (upstream or interface). You cannot configure multiple load-balance groups to distribute downstreams or upstreams that share physical connectivity.

All upstream ports coming from the same splitter must use different center frequencies that are separated by the channel width. For example, if the upstreams are using a channel width of 3.2 MHz, the center frequencies for all upstreams must be separated by at least 3.2 MHz.

All downstreams and upstreams in a load-balance group must share physical connectivity to the same group of cable modems. Downstreams can be in a separate load-balance group than upstreams, but all downstreams or all upstreams that have the same RF physical connectivity must be members of the same load-balance group.

Load balancing is not supported across cable linecards. Load balancing between interfaces on the same linecard is allowed.

Shared spectrum groups are not supported when using load balancing across multiple interfaces that support spectrum hopping (MC28U, MC5x20U). When load balancing is configured across these interfaces, you must configure non-overlapping and separate spectrum groups on each interface. If you have configured upstream shared spectrum groups while doing downstream load balancing, the downstream in each MAC domain must not use overlapping upstream groups. For example, the downstream in one MAC domain could use an upstream spectrum band of 10 to 30 MHz, while the downstream in a second MAC domain could use an upstream spectrum band of 30 to 42 MHz. Each MAC domain has its own upstream shared spectrum group, allowing the load-balance group to contain the downstreams for both MAC domains. These bands must also be big enough to supply the bandwidth needed for all the USs that share the band plus enough for hopping.

Do not use the utilization method of load balancing on cable interfaces that have a small number of cable modems and where a single modem is responsible for the majority of the interface load. In this condition, the Cisco CMTS could end up continually moving cable modems from one interface to another in an endless attempt to load balance the interfaces. In the situation where one cable modem creates the bulk of the load on an interface, the load-balancing thresholds should be configured for a value above the load created by that single modem.

In Cisco IOS Release 12.2(15)BC1, the dynamic load balancing method uses the Downstream Frequency Override message to move cable modems between downstream channels, which results in cable modems going offline and having to reregister, resulting in a short, temporary loss of connectivity for the customer. This is because the DOCSIS 1.0 specification requires cable modems to reregister whenever the downstream is changed using the Downstream Frequency Override message. Cable modems should not go offline when they are moved between upstreams.

It is highly recommended to use interface bundling on all interfaces that belong to the same load-balance group. Interfaces in the same load balance group must be in the same interface bundle. If interface bundling is not incorporated and DS members of the same load-balance group are in multiple IP domains, this will cause hosts to lose connectivity due to changed subnets.

If you have configured load balancing, the provisioning system must not assign specific upstream channels or downstream frequencies to individual cable modems in their DOCSIS configuration files. If specific upstream or downstream frequencies are provisioned, then those modems must be excluded from load balancing.

Load balance is supported on all linecards in the uBR 10K, uBR7246VXR and uBR7100 with 12.3 (9) IOS and later.

### **Example: Method Modem -** Default method

**Note:** 1. The Load-balancing algorithm uses a logical "**AND**" function where the minimum threshold and the load threshold are used to determine if static load-balancing should be turned on. In dynamic load-balancing, the enforce threshold and the minimum threshold are used to determine if dynamic load-balancing should be turned on.

The load-balancing algorithm, when turned on, will attempt to balance to an absolute zero difference in load between all ports in that load-balance group. Keep in mind that this may appear different than expected due to port weight calculations. Port weight is directly related to the raw throughput of that port. The load-balancing algorithm will continue to balance until either an absolute zero difference in load exists or when checked during the next interval and any one of the thresholds are met. If you have 300 modems to balance and 180 are on US0 and 120 are on US 1, then the % difference is (300/2 - 120)/300 = (150-120)/300 = 30/300 = 10%.

The following is a description of commands used with load balancing when configuring the modem method (default).
Router(config)#cable load-balance group 1 interval 10
Router(config)#cable load-balance group 1 threshold load minimum 2
Router(config)#cable load-balance group 1 threshold load 10 enforce 15

## Basic Wiring for Upstream Load Balancing

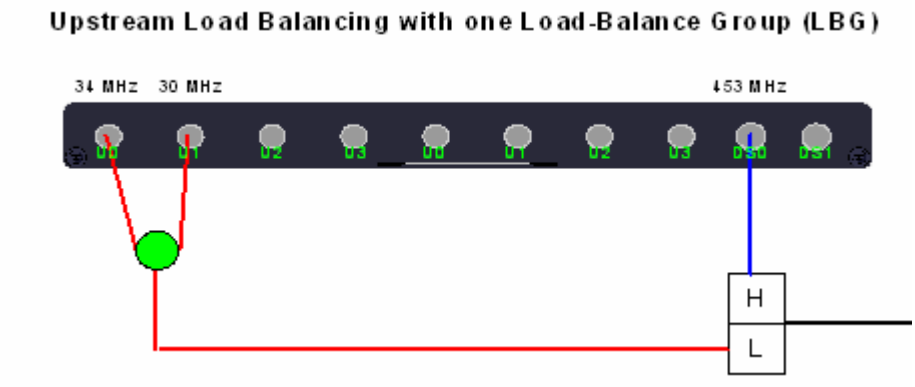


Figure 1

Load-Balance Method: modem (*static is the default setting)

Router(config)#cable load-balance group 1
Router(config)#int cx/y
Router(config-if)#cable u0 load-balance group 1
Router(config-if)#cable u1 load-balance group 1
!* Default settings for load-balancing are static; “registration-based”

Load-Balance Method: modem (**dynamic)

Router(config)#cable load-balance group 1 threshold load minimum 2
Router(config)#cable load-balance group 1 threshold load 10 enforce 15
!** Dynamic settings allow user configurable thresholds for “post-registration” load-balancing

Router(config)#int cx/y
Router(config-if)#cable u0 load-balance group 1
Router(config-if)#cable u1 load-balance group 1

## Entire MAC Domain Load Balancing

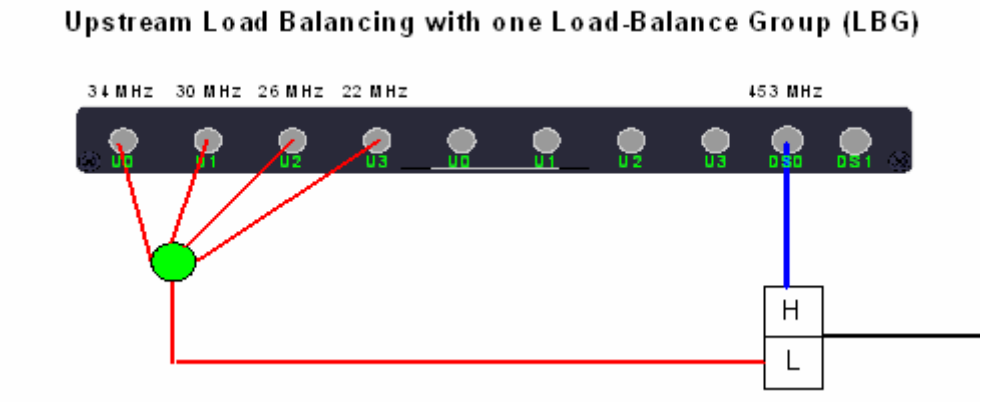


Figure 2

Load-Balance Method: modem (static)

Router(config)#cable load-balance group 1
Router(config)#int cx/y
Router(config-if)#cable u0 load-balance group 1
Router(config-if)#cable u1 load-balance group 1
Router(config-if)#cable u2 load-balance group 1
Router(config-if)#cable u3 load-balance group 1

Load-Balance Method: modem (*dynamic)

Router(config)#cable load-balance group 1 threshold load minimum 2
Router(config)#cable load-balance group 1 threshold load 10 enforce 15
!* Dynamic settings allow user configurable thresholds for “post-registration” load-balancing

Router(config)#int cx/y
Router(config-if)#cable u0 load-balance group 1
Router(config-if)#cable u1 load-balance group 1
Router(config-if)#cable u2 load-balance group 1
Router(config-if)#cable u3 load-balance group 1

## US & DS Load Balancing

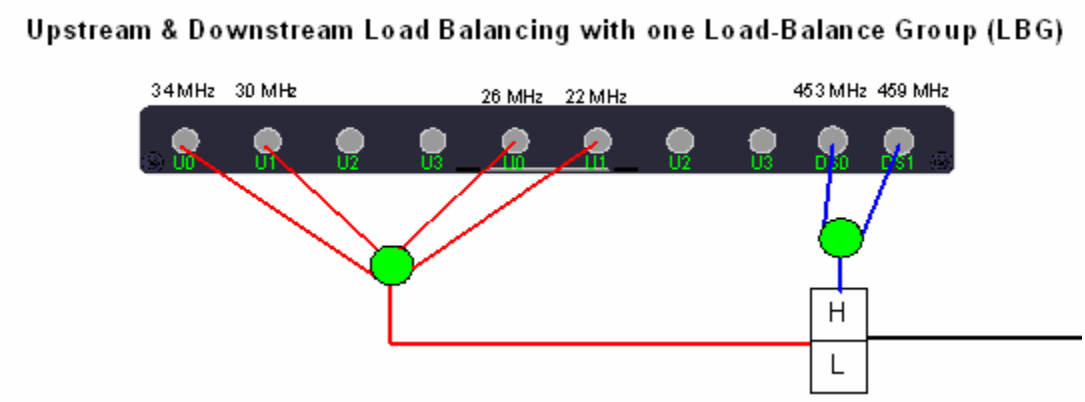


Figure 3

Load-Balance Method: modem (static)

Router(config)#cable load-balance group 1
Router(config)#int cx/y
Router(config-if)#cable downstream frequency 453000000
Router(config-if)#cable load-balance group 1
Router(config-if)#cable u0 load-balance group 1
Router(config-if)#cable u1 load-balance group 1
Router(config)#int cable a/b
Router(config-if)#cable downstream frequency 459000000
Router(config-if)#cable load-balance group 1
Router(config-if)#cable u0 load-balance group 1
Router(config-if)#cable u1 load-balance group 1

Load-Balance Method: modem (*dynamic)

Router(config)#cable load-balance group 1 threshold load minimum 2
Router(config)#cable load-balance group 1 threshold load 10 enforce 15
!* Dynamic settings allow user configurable thresholds for “post-registration” load-balancing

Router(config)#int cable x/y
Router(config-if)#cable downstream frequency 453000000
Router(config-if)#cable load-balance group 1
Router(config-if)#cable u0 load-balance group 1
Router(config-if)#cable u1 load-balance group 1
Router(config)#int cable a/b
Router(config-if)#cable downstream frequency 459000000
Router(config-if)#cable load-balance group 1
Router(config-if)#cable u0 load-balance group 1
Router(config-if)#cable u1 load-balance group 1

## Example: Method Utilization

The following is a description of commands used when configuring the utilization method.
Router(config)#cable load-balance group 1 method utilization
Router(config)#cable load-balance group 1 interval 10
Router(config)#cable load-balance group 1 threshold load 10 enforce 15

**Note:** the "load minimum x" command is not used for utilization method.

The load balancing algorithm uses a function where the load threshold is used to determine if static load balancing should be turned on. In dynamic load balancing, the enforce threshold is used to determine if dynamic load balancing should be turned on.

The load balancing algorithm, when turned on, will attempt to balance to an absolute zero difference in load between all ports in that load-balance group. The load balancing algorithm will continue to balance until either an absolute zero difference in load exists or when checked during the next interval and any one of the thresholds are met.

**Note**: If the total utilization is less than 25% then no load balancing will be done when using the utilization method.

Utilization method is not recommended for static (registration-based) load balancing. This is because a single modem using a substantial percentage of one upstream's bandwidth will cause modems to be unevenly distributed. When that one modem is no longer sending traffic, there will be a lopsided number of modems distributed within the load balance group. Since it is setup for static load balancing, the modems will stay in this lopsided state until more modems come online (which could require hundreds) or there is an event which resets the interface (shut/no shut, router reboot,..)

Load balancing uses a long term average (in the latest code release), such that the US utilization calculation matches the DS calculation. This was done after a bug was filed against it. The mac-scheduler command displays a short term average (every 1/2 second).

## Basic Wiring for Upstream Load Balancing

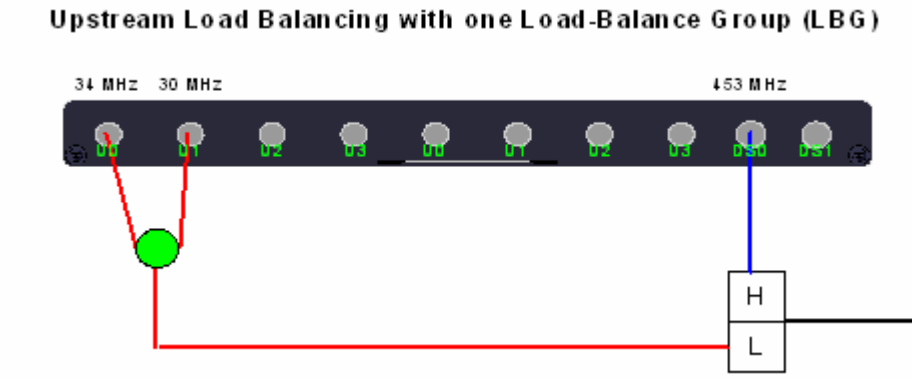


Figure 4

Load-Balance Method: utilization (static)

Router(config)#cable load-balance group 1 method utilization
Router(config)#cable load-balance group 1
Router(config)#int cx/y
Router(config-if)#cable u0 load-balance group 1
Router(config-if)#cable u1 load-balance group 1

Load-Balance Method: utilization (*dynamic)

Router(config)#cable load-balance group 1 method utilization
Router(config)#cable load-balance group 1 threshold load 10 enforce 15
!* Dynamic settings allow user configurable thresholds for “post-registration” load-balancing.

Router(config)#int cx/y
Router(config-if)#cable u0 load-balance group 1
Router(config-if)#cable u1 load-balance group 1

## Entire MAC Domain Load Balancing

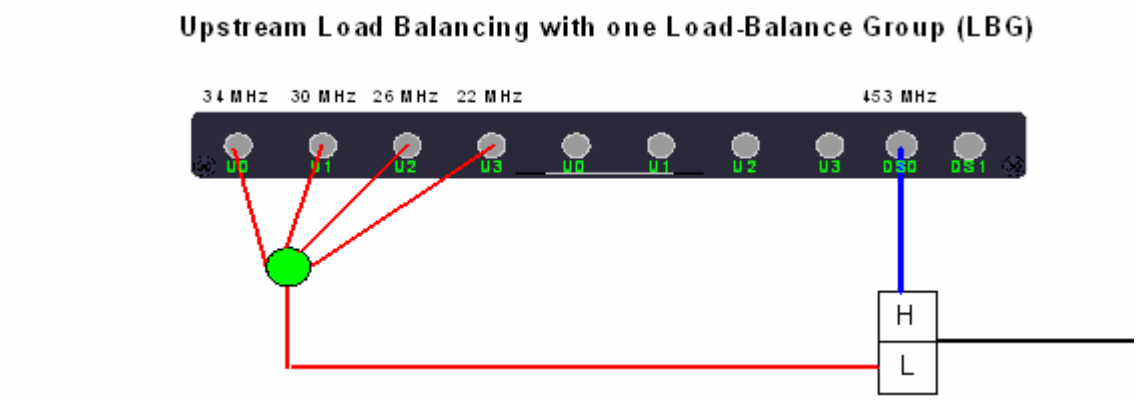


Figure 5

Load-Balance Method: utilization (static)

Router(config)#cable load-balance group 1 method utilization
Router(config)#int cx/y
Router(config-if)#cable u0 load-balance group 1
Router(config-if)#cable u1 load-balance group 1
Router(config-if)#cable u2 load-balance group 1
Router(config-if)#cable u3 load-balance group 1

Load-Balance Method: utilization (*dynamic)

Router(config)#cable load-balance group 1 method utilization
Router(config)#cable load-balance group 1 threshold load 10 enforce 15
!* Dynamic settings allow user configurable thresholds for “post-registration” load-balancing.

Router(config)#int cx/y
Router(config-if)#cable u0 load-balance group 1
Router(config-if)#cable u1 load-balance group 1
Router(config-if)#cable u2 load-balance group 1
Router(config-if)#cable u3 load-balance group 1

## US & DS Load Balancing

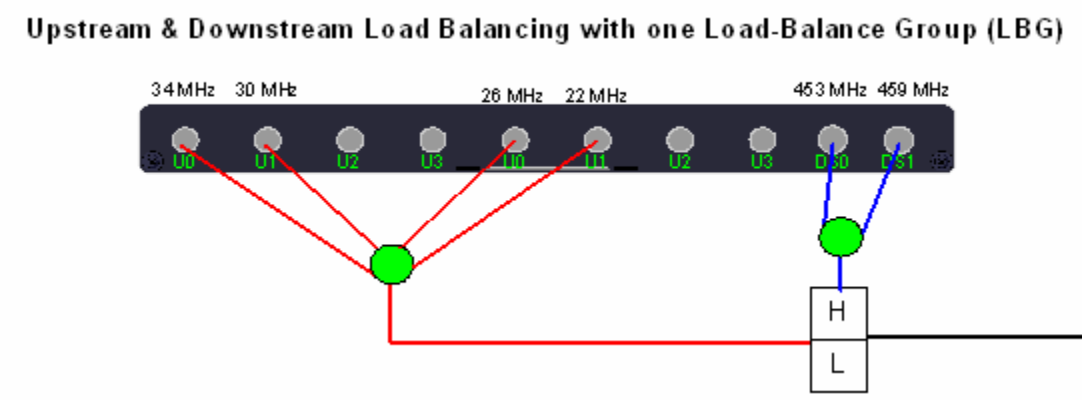


Figure 6

Load-Balance Method: utilization (static)

Router(config)#cable load-balance group 1 method utilization
Router(config)#int cx/y
Router(config-if)#cable downstream frequency 453000000
Router(config-if)#cable load-balance group 1
Router(config-if)#cable u0 load-balance group 1
Router(config-if)#cable u1 load-balance group 1
Router(config)#int cable a/b
Router(config-if)#cable downstream frequency 459000000
Router(config-if)#cable load-balance group 1
Router(config-if)#cable u0 load-balance group 1
Router(config-if)#cable u1 load-balance group 1

Load-Balance Method: utilization (*dynamic)

Router(config)#cable load-balance group 1 method utilization
Router(config)#cable load-balance group 1 threshold load 10 enforce 15
!* Dynamic settings allow user configurable thresholds for “post-registration” load-balancing.

Router(config)#int cable x/y
Router(config-if)#cable downstream frequency 453000000
Router(config-if)#cable load-balance group 1
Router(config-if)#cable u0 load-balance group 1
Router(config-if)#cable u1 load-balance group 1
Router(config)#int cable a/b

```
Router(config-if)#cable downstream frequency 459000000
Router(config-if)#cable load-balance group 1
Router(config-if)#cable u0 load-balance group 1
Router(config-if)#cable u1 load-balance group 1
```

## Exclude Command - Details

One can exclude a particular CM, or all CMs from a particular vendor from one or more forms of load-balancing operations.

`(no) cable load-balance exclude {modem` *`mac-address`* `| oui` *`oui-value`*`} [enforce | static | strict]`

- **modem *mac-address*** - Specifies the hardware (MAC) address of an individual cable modem to be excluded from load balancing. (You cannot specify a multicast MAC address.)
- **oui *oui-value*** - Specifies the organization unique identifier (OUI) of a vendor, so that cable modems from this vendor are excluded from load balancing. The OUI must be specified as three hexadecimal bytes separated by either periods or colons.
- **enforce** - (Optional) Excludes the cable modems from dynamic load balancing, but they continue to participate in static load balancing.
- **static** - (Optional) Excludes the cable modems from normal static load balancing. They continue to participate using a passive type of static load balancing and dynamic load balancing.
- **strict** - (Optional) Excludes the cable modems from all forms of load balancing.
- **static strict** - (Optional) Excludes the cable modem from all types of static load balancing, but they continue to participate in dynamic load balancing.

| | Static | Passive | Dynamic |
|---|---|---|---|
| **(no cable load-balance exclude) Default** | **Not Excluded** | **n.a** | **Not Excluded** |
| **cable load-balance exclude** | **Excluded** | **Not Excluded** | **Excluded** |
| **cable load-balance exclude enforce** | **Not Excluded** | **n.a** | **Excluded** |
| **cable load-balance exclude static** | **Excluded** | **Not Excluded** | **Not Excluded** |
| **cable load-balance exclude strict** | **Excluded** | **Excluded** | **Excluded** |
| **cable load-balance exclude static strict** | **Excluded** | **Excluded** | **Not Excluded** |

**Table 1**

Examples:

```
cable load-balance exclude oui 00.04.00
cable load-balance exclude oui 00.03.00 static
cable load-balance exclude oui 0C.00.00
cable load-balance exclude modem 0001.0203.0405 static
cable load-balance exclude modem 0C0B.0A09.0807
```

## Load Balance Test Command

A test command exists to move modems while observing a debug-type of output.

```
Router# test cable load-balance 0000.394e.4e59
Sending UCC request: Cable5/0/0/U0 --> U1
          Waiting for test completion ........
          Test results:
          UCC Response: 0.0s
          Initial Ranging: 8.5s
          Ranging Complete:
     Modem replied to DOCSIS ping.
     Test summary:
     UCC Response: success rate 100% min 0.0s max 0.0s avg 0.0s
     Initial Ranging: success rate 100% min 8.5s max 8.5s avg 8.5s
     Testing US Channel Override: Cable5/0/0/U1 > U0
          Waiting for test completion ........
     Test results: Initial Ranging: 8.5s
     Ranging Complete: failed.
     Modem replied to DOCSIS ping.
     Test summary:
     UCC Response: success rate 100% min 0.0s max 0.0s avg 0.0s
     Initial Ranging: success rate 100% min 8.5s max 8.5s avg 8.5s
```

## Application Note: Utilization

The “Exclude Modem” Load-Balancing command can be used as a mechanism for allowing unused bandwidth to be used by residential customers on US ports that have been configured for business use but are under utilized. When Business users begin to saturate their US port, the residential users can be moved off that US port.

Example:
1. USs 0,1,2,3 are all part of same US LB group.
2. US 0 has business customers with DOCSIS config files assigning it to US 0, specifically.
3. A LB group is configured for method utilization.
4. Use the “exclude modem” CLI command to prevent the business customer modem from ever being moved to another US port.
5. The residential customers have DOCSIS config files that do not contain US channel information

As a result of the above configuration, when bandwidth utilization on US 0 (primarily business customers) drops to a point where the difference exceeds the “enforce” threshold, residential

customer modems could be moved to the under utilized US 0 until the business customer needs the bandwidth. At this time the residential customer modems would be moved to another US port. The only problem is, what if the other US ports are saturated also?

## MTA Load Balance Example:

A customer does load balancing between 2 US ports based on modem count and allows 200 modems per US. Some commercial or home office subscribers have a CIR of 128 kbps. By configuring the original command of cab upstream 0 admission control 40%, only 10.24 * .4 / .128 = 32 modems would be allowed to register per port with this config file. They expect a 60% take rate on G711, 20 msec phone service leading to 200 * .6 = 120 MTAs with a 1:6 oversell equaling 20 simultaneous calls.

**Note:** Admission control for UGS flows is a feature slated for CY2005. Understanding upstream (US) UGS rates and how admission control will be calculated helps assure proper results are achieved. More information can be found listed under DDTS, CSCeg24047 AC10: US bit rate reporting should account for scheduler inefficiency and EDCS-416700 (Upstream (US) Overhead Assumptions for VoIP).

The current admission control feature could be used to control the amount of MTAs per US port, but it would only be relevant for registration-based load balancing. All the modems that have MTAs would need to be assigned a committed information rate (CIR).

To guarantee this 1:6 oversell assumption, we would configure a critical threshold of .1152 * 20 / 8.192 = **28%**. The overall affect would be 115.2 * 20 = 2.3 Mbps used for UGS when active.

Interestingly, the CIR modems with a minimum of 128 kbps would need to include ~ 20% for FEC and physical layer overhead. This would lead to 32 * .128 * (1+.2) = 4.91 Mbps for CIR if used, and the rest would be leftover for best effort. The leftover would be 8.192 - 2.3 – 4.91 = .98 Mbps. So, this would be leftover for the 200 – 20 - 32 = 148 modems. In reality, the maintenance overhead would be slightly less than the 20% assumed from the 10.24 Mbps Raw rate giving slightly more throughput for the best effort modems. Most assumptions for best effort would conclude that users would achieve ~ (8.7-2.3-4.9) / (148*.10) = 101 kbps.

## Bad Examples and FAQs

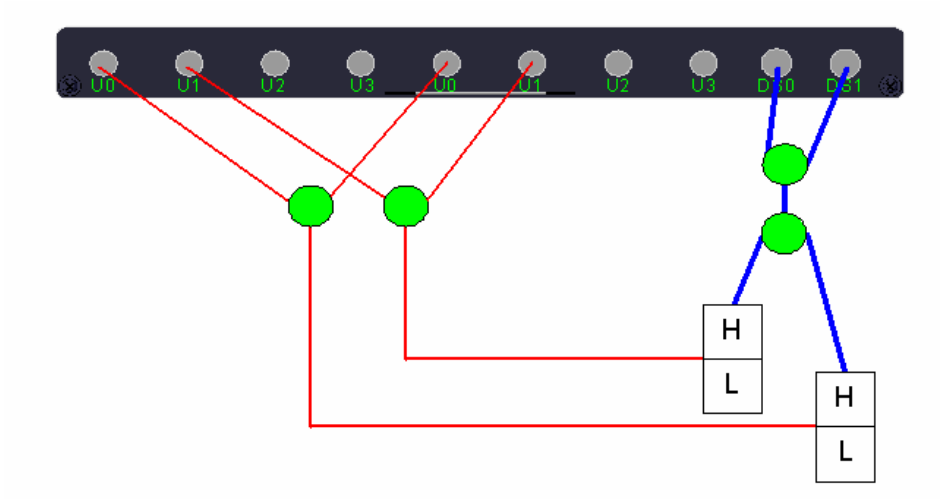


Figure 7

This example would not allow proper load balancing. USs that share a common splitter feed 2 mac domains. Example, 200 CMs total with 50 per US (50 50 50 50).

If US 0 dies, you could end up with 100 on US 1 because the DSs stay balanced & the result looks like (0 100 100 0). A new feature is planned that will allow pure US load balance without DS concerns so it would look like (0 50 100 50).

Q. Can DS load-balance group 1 be re-utilized in another MAC Domain?
A. It can not be re-used across all DSs because the LB code will attempt to load balance across those DSs even though they don't have physical connectivity.

Q. Can US load-balance group 2 be re-used in another MAC domain where one of the modulation-profiles is different?
A. If the DSs aren't in a LBG, then US LBGs can be re-used. The modulation or channel width can be different, but the load weights will be different, so the load may not look balanced even though it is.

Q. Can load-balance group 2 be re-used between upstreams in the same interface?
A. Not in the same mac domain.

## Load Balancing Going Forward

Due to Cable Labs® decision to make Load Balancing a specified and documented feature, Cisco's continuing efforts at improving and evolving the feature have halted. Cable Labs® is expected to define algorithms and API specifications for CMTSs at which time the feature will be revisited. Please see CSCeb01694, which is listed below:

HELD STATE INFORMATION Held awaiting CableLabs
********** Description (Modified 040326 by groeck) **********
This ddts entry is to track future (post-FCS) load balancing enhancements. Items listed below are to be considered ideas. Individual items may or may not be implemented.

1) Use DCC for DOCSIS 1.1 modems to move registered modems from one interface to another, and from one upstream to another without re-registering.

2) Create a database of "rogue" OUIs. This database would list modem responses to move requests (interface, upstream, UCC, and DCC). The system would then use this database to determine how to move modems with the same OUI. For example, if "US Override" (or the chosen default method) fails for a specific set of modems, the system would use "passive" mode for this set of modems. TBD would be the "triggering" failure rate (e.g., 50%). This database could be kept on a per-modem basis, but such an approach would not scale well and have a slow response time, especially during bring up. A dual per-OUI and per-modem database might be considered. One might consider keeping success/failure statistics with the existing modem database (modem instance). This would resolve the scaling problems. To cover changes due to modem code reload, one might consider less frequent retries for affected modems. Also, the failure data should be cleared with "clear cable modem state".

3) Complete N+1 support. After the ongoing N+1 CLI rework is complete, need to re-validate load balancing support in conjunction with N+1 and implement necessary fixes/improvements.

4) For ubr10k, remove active part of centralized code (sending modem activation data to all cards). Instead, implement a de-centralized mechanism. Note Done. Still need to optimize (reduce) number of IPC messages sent from card to card.

5) Fix downstream channel ID issues. Set unique and well defined DS channel ID (default) per chassis. Use load balancing mechanism to distribute channel IDs across cards. Send DS channel override based on this data.

6) Implement connectivity statistics on a per-connection basis (e.g. Petri network). Possibly include "auto" mode where nodes may be added dynamically based on gathered information (e.g. modems showing up on different interfaces over time. Specifically, channels not assigned to a load balancing group would be dynamically added if the system determines that connectivity exists.

7) Implement test command to test downstream/interface load balancing. Currently implemented;

```
R7636-ubr10k#test cable load 00a0.73b0.4bdd ?
  <1-20>    Test repeat count
  ucc        Upstream Channel Change
  upstream   Upstream channel override
```

We need the same for downstream.

8) Intelligence needed when moving modems with CIR (QoS). In the present design, the selection of modem to be moved (to balance the load) is random from the available and allowed modems to be moved. Intelligence should be added to either detect a modem with CIR (QoS) and not move it if the new US does not have BW to service it, OR it should be configurable via CLI (like UGS) to stop (exclude) the CIR modem or other type of QoS (if applicable) modem from moving. When utilization is used as the method this may not be a problem, as modems are moved from higher utilization to lower. But when using modem or SF as the method we can run into this scenario.

9) System should be able to automatically assign US channels to a LB group if movement of modems is detected. This would be very helpful if the user sets an interface LB group but does not

assign US channels to groups. For "new" groups might use internal numbers which can not be assigned by users. Maybe similarly can be done if movement from one interface to another is detected.

10) Provide configuration option to apply interface load balancing to associated upstream channels only. In other words, for such interfaces, no attempt would be made to balance the interface load, but the upstream channels in the same group would get perfectly balanced, even across interfaces.

11) Define minimum threshold to apply to utilization (besides # of modems).

12) [Depending on CableLabs mandate] Add capability to override group thresholds on a per channel basis.

13) Do not load balance modems with an upstream channel ID in the configuration file.

14) Add support for time based dynamic load balancing, i.e., only actively move modems during a limited period of time.

15) Add a run time command to load balance dynamically even if the configuration is for static. This is to allow the system to be load balanced the very first time. Right now the user has to configure dynamic, then change the config back to static if that is the actual intent.

16) If method is utilization, add a CLI flag which enables modem based load balancing after interface initialization, until the minimum thresholds for utilization based load balancing is reached; this flag (and modem based load balancing) would then be turned off and stay off until the load balancing state moves back to "init“.

## References


- **Configuring Load Balancing for the Cisco CMTS**
- **DOCSIS 1.1 RFI Spec - SP-RFIv1.1-I09-020830**
- **Load Balance power point slides - EDCS-313313**


# Appendix

## Legacy Load Balancing on the Upstream

These are some ideas to achieve Upstream (US) load balancing in regards to the number of modems per US port. Before covering the different suggestions, keep in mind the following information:

1. Using multiple ports requires multiple US frequencies
2. More insertion loss from a splitter
3. Noise is still funneled from all the customers
4. Try to keep the modems per US port to no more than 150 to 200 per US, but this is subjective. If doing VoIP, you may want to cut this in half. However, advances in DOCSIS Phy technology may allow greater US aggregated bandwidth allowing more modems per US than currently recommended. Devices such as a digital settops requiring low bandwidth may also be installed and allow more devices to be installed. For guidelines to the

maximum number of recommended users on an upstream or downstream port, refer to What is the Maximum Number of Users per CMTS document.

5. Separate the center frequencies by 3.2 MHz for future channel width upgrades.
6. If using an S-card and assigning spectrum groups, make spectrum bands 3.2 MHz wide and separate by at least 20 kHz. Also assign them as "shared spectrum". This allows N+1 redundancy for frequency hopping if an extra band is assigned.

The following 4 suggestions are possible for upstream load balancing with legacy IOS.

1. Setup US 0 for 1 freq. and US 2 for another freq. Wire one US node Rx to both US ports, shut/no shut those 2 US ports, and then pray ☺. You could bounce the whole interface, but that would take down all the modems on the DS mac domain. Cisco sends the UCDs in a pseudo-random fashion on newer IOS so the modems "should" balance out between the 2 USs. This should work ok when new modems are coming online, but may not work that well for modems that have already cached an existing UCD.

The following table lists the pseudo-random order of the UCDs.

Upstream Allocation Sequence

| Timeslot | 1st choice | 2nd | 3rd | 4th | 5th | 6th |
|---|---|---|---|---|---|---|
| A | 0 | 1 | 2 | 3 | 4 | 5 |
| B | 5 | 0 | 1 | 2 | 3 | 4 |
| C | 4 | 5 | 0 | 1 | 2 | 3 |
| D | 3 | 4 | 5 | 0 | 1 | 2 |
| E | 2 | 3 | 4 | 5 | 0 | 1 |
| F | 1 | 2 | 3 | 4 | 5 | 0 |

Knowing the sequence of UCDs, can help to determine the best way to physically combine US ports. If 3 US ports will be combined, combine even ports together (0, 2, & 4) and odd ports together (1, 3, & 5). This will result in a distribution that looks like this:

| | HFC Area 'A' | | | HFC Area 'B' | | |
|---|---|---|---|---|---|---|
| Timeslot | US0 | US2 | US4 | US1 | US3 | US5 |
| A | X | | | X | | |
| B | X | | | | | X |
| C | | | X | | | X |
| D | | | X | | X | |
| E | | X | | | X | |
| F | | X | | X | | |
| Count | 2 | 2 | 2 | 2 | 2 | 2 |

If only 2 ports are used, combine (0 3), (1 4), and (2 5) for perfect balance.

2. If the modems are already scattered between the 2 USs, you can force specific modems over to a certain US port by giving some modems a DOCSIS config file with the US channel ID of 1 for US 0 and other modems a config file with an US channel ID of 2 for US 1, but the modems must load the new config file, which also has to be developed. Channel ID 0 is reserved for Telco return. **Caveat**: What happens if you drop an US port because of noise at that freq. or a hardware problem, the modems that are hard-coded to use that US will never register on the other.

3. If the modems are already scattered between the 2 USs, you can force specific modems over to a certain US port without shutting the interface or ports. Use the test command; **test cable ucc cx/y {sid #} {Port #}**. The modem is supposed to change US ports without rebooting. This could take you a while to do each one individually, unless you write some type of prl script. This command is also not supported such as other test commands. You can also specify the DOCSIS 1.1 ranging technique to use when moving the cable modem:
   0 = Specifies that the modem should perform initial maintenance on the new channel.
   1 = Specifies that the modem should perform only station maintenance on the new channel.
   2 = Specifies that the modem should perform either initial maintenance or station maintenance on the new channel.
   3 = Specifies that the modem should immediately use the new channel without performing initial or station maintenance.

   You can also do, **`cable modem {mac/IP-address} change-frequency {Ch ID #}`**. Ch 1 = US 0 Ch 2 = US 1, etc. The problem with this command is that it will force the modem to re-acquire and thus, go offline first.

```
ubr10k#cable modem 0003.e3fa.5e11 ?
  change-frequency   Change DS frequency or US channel id
  max-hosts          Modem max hosts number
  qos                Cable quality of service

ubr10k#cable modem 0003.e3fa.5e11 change-frequency ?
  <1-6>                    Upstream Channel ID
  <54000000-1000000000>    Downstream Frequency in Hz
```

   **Note**: The following command allows you to force a different QoS profile to a certain modem, but it only works for DOCSIS 1.0 modems.

```
ubr10k#cab modem “mac-address” qos profile ?
<1-255>  Qos class index
```

   This command may be present, but not activated in current DOCSIS 1.1 IOS (12.2(15)BC1).

4. Setup US 0 for 1 freq. and US 1 for another freq, wire one US node Rx to both US ports, make DOCSIS config files with a minimum guarantee rate, assign “Admission Control” on the US ports of 100%, shut/no shut those 2 US ports. See: http://cable.cisco.com/documents/AdmissionControl.html.

   Example: 400 modems total between 2 US ports, running QPSK a 1.6 MHz wide channel. The symbol rate will be 1.28 Msym/s with 2 bits/sym = 2.56 Mbps total aggregate throughput. If I want 200 modems on each US at 100% admission control, that would be 2.56 Mbps/200 = 12.8 kbps. The Min Guarantee US rate would be set for 12.8 kbps. This means that once I reach 100% of 2.56 Mbps, no more modems will be allowed to register on that particular US and will try another UCD.

I agree that **modem** balancing does not constitute **true load** balancing. That would require looking at a command like admission control, but enforcing it after the modem has registered. The admission control command only enforces the registration process of the modems. A new load balance command is available now that allows modems to be balanced based on utilization, service

flows, or modem count. Another idea is balancing based on different types of traffic to be weighted. If you only want 30% UGS traffic, that could be administered.

**Tip**: Use the following command to see how much bandwidth has been allocated:

```
ubr10k#show interfaces c5/0/0 mac-scheduler 0 | inc SIDs
     UGS    : 0 SIDs, Reservation-level in bps 0
     UGS-AD : 0 SIDs, Reservation-level in bps 0
     RTPS   : 0 SIDs, Reservation-level in bps 0
     NRTPS  : 0 SIDs, Reservation-level in bps 0
     BE     : 2 SIDs, Reservation-level in bps 1024000
```

In this example, each modem was guaranteed a min of 512 kbps and only 2 modems are online.

With the release of 12.3 (9a) IOS, you can allocate a percentage of throughput to certain flow types. The command is:

**cable admission-control {scheduling-type** *flow-type* | **service-class** *name*}
*reservation-limit*

The following command indicates how the allocation was administered and the current usage.

```
ubr10k#show cable admission-control
Service      Limit         Current       Scheduling     Limit         Current
Class        (% of BW)     Usage (%)     Type           (% of BW)     Usage (%)
----          0              0            BE             30           30
----          0              0            NRTPS          X             X
----          0              0            RTPS           20           20
----          0              0            AD-UGS         30           27
----          0              0            UGS            20           13
```

This is for the whole chassis.