



## Traffic Forecasting

---

WAE Design traffic forecasting tools let you create plans containing estimates of future traffic based on projected growth of current measured or simulated traffic. From these new plans you can determine the impact of the new traffic on the network and plan for upgrades. The traffic growth estimates used in this process can be manually entered, for example, based on knowledge of new services being introduced. After creating these growth plans, you can analyze their capacity requirements using Simulation Analysis, and use the GUI to upgrade the network to meet these requirements.

One example use of these growth plan tools would be to predict the effects of a certain percentage of increased traffic for each quarter of the coming year. You could generate a separate plan for each of these traffic predictions to forecast the impact of each quarter's anticipated growth. Another example application would be to group the demands into demand groupings to predict the effects of different traffic growth at each site, such as 10% for Paris, 20% for London, and 30% for Amsterdam.

## Demand Groupings

Demand Groupings define a group of demands. They provide a convenient way of specifying aggregated traffic in a plan, which can be used as a basis for traffic reports and growth plans. For example, you can use demand groupings to represent the following.

- Total traffic sourced from one specific site
- Total VPN traffic, or any defined service class, sourced from one specific site to another
- Total traffic destined for a particular AS

To generate a report on demand groupings, including costs associated with them, choose **Tools > Reports > Demand Groupings**. For detailed information, see [Demand Groupings Reports](#). For a description of the cost modeling columns in the report output, see [Cost Modeling](#).

## Demand Groupings and Their Filters

The term *demand grouping* refers to a set of filters (rules) that determine a group of demands. For example, you could create a demand grouping that filters all demands sourced from Paris, and name it ParisSrc. When you apply the filter, if 52 of the existing demands contain Paris as a source, those 52 demands combined comprise a single demand grouping called ParisSrc.

Each row in the Demand Groupings table identifies a uniquely named demand grouping and the associated filters that define which demands are in it.

- Each column other than Name is a criterion that a demand must meet to be in that demand grouping.

- A demand must meet all criteria defined in these columns in order for it to be a member of the demand grouping.
- If the field is empty, it is not used as a criterion.
- The column headings with “Matches” in the name use regular expressions and the filters are applied accordingly.

Example: If Destination Node Matches is `*\..atl$`, a demand can be a member of this demand grouping only if it has destination nodes ending in `.atl` and if it meets all other criteria (listed in the other columns).

- The column heading that use “Equals” in the name must be exact matches.

Example: If Source Site Equals is `chi`, a demand can be a member of this demand grouping if its source site is named exactly `chi` and if it meets all other criteria.

- If the Tags Include column has an entry, a demand must contain this tag to be included in this demand grouping. (Demands can have multiple tags.)

Example: If Tags Include is `atlantic`, and if a demand has three tags named `europa`, `atlantic`, and `apac`, then the demand would be included in this demand grouping provided it meets all other criteria.

There are other object-based Tags Include columns for source and destination sites, nodes, and ASes. If one of these columns has an entry, the source or destination object must contain this tag to be included in this demand grouping.

Example: If Source Site Tags Include is `Europe` and if the Dest Site Tags Include is `Asia`, the demand would be included in this demand grouping if its source site is tagged with `Europe`, if its destination site is tagged with `Asia`, and if it meets all other criteria.

- Demand Grouping columns can contain information that does not exist in the current plan. This flexibility allows a demand grouping to be applied to any network.

To see a list of demands that comprise a demand grouping, in the Demands Grouping table, right-click the row containing the unique demand grouping name and select Filter to Associated Demands.

- In [Figure 7-1](#), if you were to filter to associated demands on the first row (`chi` to `bos`), the resulting Demands table would show four demands because they are the only ones with both `chi` as the source and `bos` as the destination.
- If you were to filter to associated demands on the second row (`chi` to `All`), the resulting Demands table would show all demands with `chi` as the source, no matter the destination.

To see all demand groupings in which a demand exists, in the Demands table, right-click the row containing the demand of interest and select Filter to Demand Groupings.

- In [Figure 7-2](#), if you were to filter to associated demand groupings on the second row (`cr1.chi_cr1.mia`), the resulting Demand Groupings table would show the two demand groupings that contain this demand (`chi` to `All`, and `All` to `mia`).
- If you were to filter to associated demand groupings on the first row (`cr1.bos_cr1.mia`), the resulting Demand Groupings table would show the only demand grouping that contains it (`All` to `mia`).

Figure 7-1 Demand Groupings Contain Demands

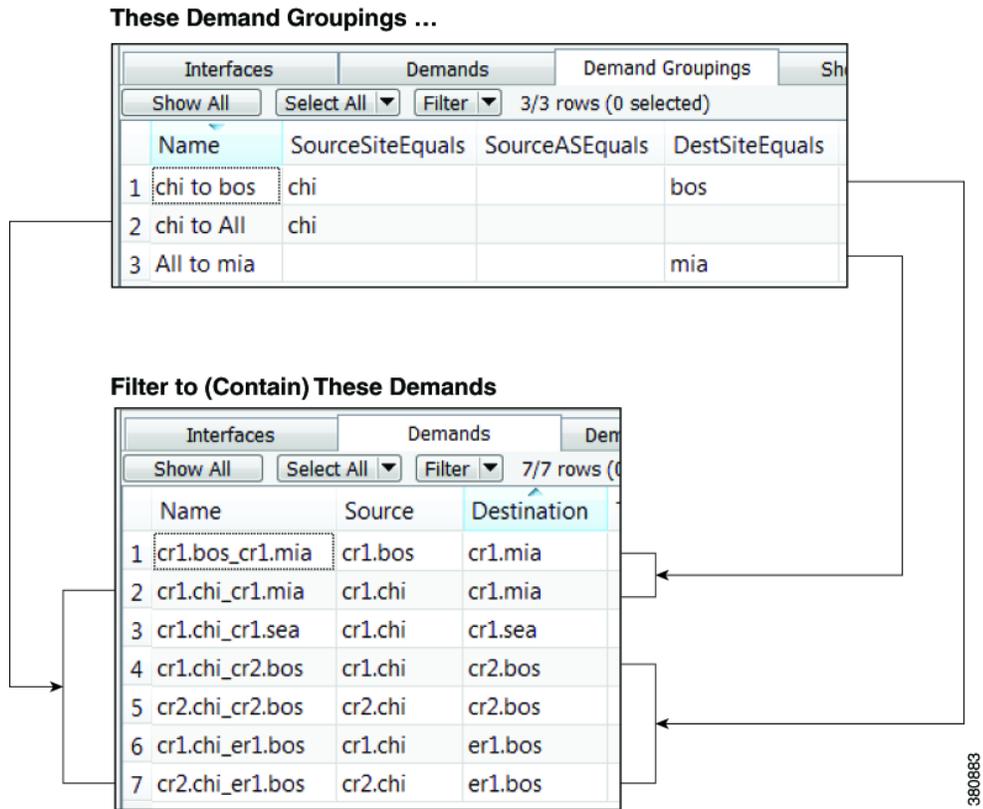
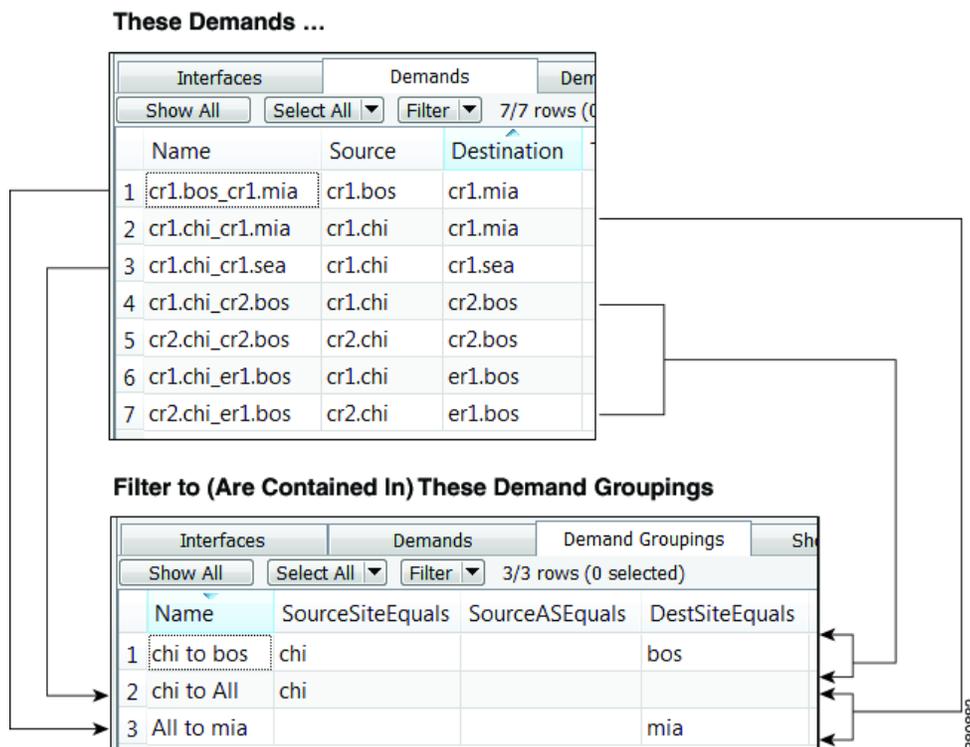


Figure 7-2 Demands Contained in Demand Groupings



## Creating a Single Demand Grouping

This method lets you create a single demand grouping. All fields other than the Name field combine to create a set of rules, or filters, for the demand grouping. A demand must meet all of the specified criteria for it to be included in the demand grouping. If you specify one or more growth values, the demand grouping can be used in creating growth plans.

Note that you can also import demand groupings. For information, see the *Cisco WAE Design Integration and Development Guide*.

- Step 1** Choose **Insert > Demands > Demand Grouping**, or right-click in the empty plot area and choose **New > Demands > Demand Grouping**.
- Step 2** Enter a unique name for the demand grouping. If you do not add a name, WAE Design renames the grouping to make it unique.

- Step 3** Enter information as needed in the Demands, Source, and Destination fields to create the set of rules that define the demand grouping. Each of these is a demand filter.
- You can enter only one parameter per field.
  - If a field is unchecked, the filter is not used.
  - If you make a selection, you must enter a value in the empty field.
    - All fields using “Tags Include” require a single tag to be specified. The list of tags for the corresponding object must include this tag.
    - All fields using “Matching” allow you to use regular expressions.
    - All fields without “Matching” require an exact match.
- Step 4** If you want to use this demand grouping when creating growth plans, enter one or more values in the Growth area. See [Single-Period Demand Grouping Traffic](#) for information on these values.
- Step 5** Click **OK**.

## Creating a Demand Grouping Mesh

A mesh of demand groupings is a set of demand groupings created between a specified set of sources and destinations. It is often more convenient to create an entire mesh at once than to create the demand groupings individually.

There are two methods for quickly creating a demand grouping mesh:

- Choose **Insert > Demands > Demand Grouping Mesh** or **New > Demands > Demand Grouping** as described the following steps.
- Importing existing demand groupings from another plan file or table file. For information on importing demand groupings, see the *Cisco WAE Design Integration and Development Guide*.

Because you are creating a set of filters, not all steps are required. By default, the Insert Demand Grouping Mesh uses selected sites (or all sites if none are selected) as the source and does not group by destination. To change the default and further restrict the demand grouping, you can choose from any combination of source, destination, service class, and demand name, as well as use object tags.

If you select multiple sites before opening the dialog box and use those sites for either the source, destination, or both, a unique demand grouping is created for each selection.

Examples: If you select both *chi* and *mia* sites for the source and nothing for destination, two demand groupings are generated: one named *chi to All* and one named *mia to All*. If you select both *bos* and *nyc* as both source and destination, four demand groupings are created: *bos to bos*, *bos to nyc*, *nyc to nyc*, and *nyc to bos*.

- 
- Step 1** Select sites, nodes, or ASes if you want to create the demand grouping for a specific set of selected objects, whether as source, destination, or both.
- Step 2** Choose **Insert > Demands > Demand Grouping Mesh**, or right-click in an empty area and choose **New > Demands > Demand Grouping Mesh**.
- Step 3** For both source and destination, you have these options:
- Sites—All, those selected prior to opening the dialog box, internal sites only (those not containing external AS nodes), or a sites with a specific tag. For example, if you selected source sites tagged with Acme, only demand groupings whose source site was tagged with Acme would be included.
  - AS's—All, those selected prior to opening the dialog box, external only, or ASes with a specific tag.
  - Nodes—All, those selected prior to opening the dialog box, internal nodes only, or nodes with a specific tag.
  - External Endpoints—All, those selected prior to opening the dialog box, or external endpoints with a specific tag.
  - Site Tags, AS Tags, Node Tags—Objects that are tagged with any value. For example, if you select “Node Tags” for destination nodes, demand groupings would include those demands whose destination nodes have one or more tags of any name. If a destination node was tagged with APAC and another destination node was tagged with Europe, both would be included in the demand grouping mesh.
  - Do not group by—Ignore sources or destinations as filters for the demand grouping meshes.

The screenshot shows a dialog box for configuring demand groupings. It has two main sections: 'Sources' and 'Destinations'. Each section contains radio buttons for different filtering criteria and dropdown menus for selecting specific objects. In the 'Sources' section, 'Sites' is selected, and the dropdown shows 'Internal (12/12)'. 'Site Tags' is also selected. In the 'Destinations' section, 'Sites' is selected, and the dropdown shows 'Internal (12/12)'. Other options like 'AS's', 'Nodes', and 'External Endpoints' are also visible with their respective counts.

- Step 4** In addition to or instead of selecting sources or destinations, you can create the grouping by restricting it to specific demands.

- Enter the exact demand name.
- Select a service class.
- Enter the exact source and destination node names or use regular expressions.

**Step 5** The demand grouping names are created based on the name or variables used in the Name of Grouping field.

- Because each demand grouping created must have a unique name, we recommend that the string “\$1 to \$2” be used as the name, which is the default. The \$1 variable is source site, node, or AS, or All if none of these are selected but you have chose to include sources in the grouping. Likewise, \$2 is the destination site, node, AS, or All if neither is selected. Thus, a name can potentially be *All to All*. If a demand grouping is selected using source and destination tags, then \$1 is the name of the source tag, and \$2 is the name of the destination tag.
- You can this field to create uniquely named demand groupings for the same site selections by entering extensions.

Example: You could select and use chi as the source, leave the destination empty, and add an extension of voice, thus creating a demand grouping named *chi to All voice*.

**Step 6** Click **OK**.

## Demand Groupings Reports

You can generate Demand Groupings reports on plan files that contain demand groupings. Each report has two sections:

- Demand Groupings—For each demand grouping in the plan file, this section lists the name, traffic level to which it applies, and the number of demands. All other columns are based on the total (aggregate) of all demands in the demand grouping, traffic, capacity bound, capacity violation, QoS bound, QoS violation, cost by capacity, and cost by utilization. These costs are the sum of the demands in the demand grouping.
  - Capacity Bound—The maximum amount that the sum of the demands in the demand grouping can be without causing over 100% utilization on any interface in the network. The calculation is based on the assumption that all demand traffic increases proportionally.

- Capacity Violation—Total traffic minus the capacity permitted for the demand grouping (Capacity Bound). A violation occurs if the Capacity Bound is exceeded. If the number appearing in the Capacity Violation column is positive, then the Capacity Bound has been surpassed. If the number is negative, the Capacity Bound has not been reached.
- For information on QoS bound and QoS violation, see [Quality of Service Simulation](#).
- For information on how these costs are calculated, see [Cost Modeling](#).
- Demands—For each demand in the plan file, regardless of whether it is in a demand grouping or not, this section lists the name, source and destination, service class, and number of demand groupings to which it belongs.

To generate a Demand Groupings report, choose **Tools > Reports > Demand Groupings**, and then click **OK**.

## Example of Capacity Bound and Capacity Violation

Each interface imposes a limit on the traffic that can be carried by all demands in the demand grouping.

The goal of this example is to determine which interface is limiting the demand grouping's ability to deliver traffic. To do this, WAE Design looks at the interfaces over which each demand is routed. WAE Design then calculates the maximum traffic that can be used for the demand without exceeding 100% utilization for any routed interface.

In this example, the demands and interfaces are as follows.

- All interfaces have 1000 Mbps of traffic.
- The demand grouping is Acme\_Region.
  - sjc\_wdc\_demand has 300 Mbps of traffic, and it traverses A\_interface and B\_interface.
  - sjc\_hst\_demand has 200 Mbps of traffic, and it traverses A\_interface and C\_interface.
- There are other demands outside this demand grouping.
- A\_interface has 600 Mbps of traffic. Of this, 500 Mbps of it is from the Acme Region demand grouping, and the remaining 100 Mbps is from demands outside this Acme\_Region demand grouping.
- B\_interface has 500 Mbps of traffic, of which 300 Mbps is from the sjc\_wdc\_demand. The remaining 200 Mbps of traffic is from demands outside this demand grouping.
- C\_interface has 400 Mbps of traffic, of which 200 Mbps is from the sjc\_hst\_demand. The remaining 200 Mbps of traffic is from demands outside this demand grouping.

The calculation is as follows:

- A\_interface
  - A\_interface can hold 900 Mbps of demand grouping traffic.  

$$\text{A\_interface capacity} - \text{non-demand grouping traffic}$$

$$1000 - 100 = 900$$
  - The amount of sjc\_wdc\_demand traffic that A\_interface can hold is 540 Mbps.  

$$\text{Total available capacity for the demand grouping} * (\text{sjc\_wdc\_demand} / \text{total demand grouping traffic})$$

$$900 * (300 / 500) = 540$$
  - The amount of sjc\_hst\_demand traffic that A\_interface can hold is 360 Mbps.

Total available capacity for the demand grouping \* (sjc\_hst\_demand/total demand grouping traffic)

$$900 * (200/500) = 360$$

- B\_interface can hold 800 Mbps of sjc\_wdc\_demand traffic.

B\_interface capacity – non-demand grouping traffic

$$1000 - 200 = 800$$

- C\_interface can hold 800 Mbps of sjc\_hst\_demand traffic.

C\_interface capacity – non-demand grouping traffic

$$1000 - 200 = 800$$

- The limiting interface is A\_interface because it imposes the strictest limits in the demand grouping's ability to deliver traffic. It supports only 360 Mbps of sjc\_hst\_demand traffic, whereas B\_interface supports 800 Mbps of sjc\_hst\_demand traffic. Also, A\_interface can support only 540 Mbps of sjc\_wdc\_demand traffic, whereas C\_interface supports 800 Mbps of sjc\_wdc\_demand traffic.

- The Capacity Bound is 900 Mbps. If the sum of the demands in the demand grouping becomes 900 Mbps, the utilization of A\_interface (the limiting interface) will be 100%. This 900 Mbps is the total available capacity for the demand grouping based on the limiting interface.

A\_interface capacity – non-demand grouping traffic

$$1000 - 100 = 900$$

- The Capacity Violation is –400.

Total A\_interface traffic permitted for the demand grouping – Capacity Bound

$$500 - 900 = -400$$

**Figure 7-3 Example Demand Groupings Report**

Reports						
Demand Groupings						
Demand Groupings (9)						
Demands (95)						
Name	Traffic Level	Traffic	Number of Demands	Cost by Capacity	Cost by Util	
sjc to All	Default	3650.60	6	58925.16	187860.95	
sjc to All	Peak	14725.30	6	235547.56	369908.59	
sjc to All	Weekend	9695.55	6	149069.67	344901.80	
All to lax	Default	5924.67	5	137336.14	451218.90	
All to lax	Peak	9290.06	5	210953.77	354144.69	
All to lax	Weekend	9453.31	5	213061.40	450689.63	
All to chi	Default	7105.47	4	97903.69	279012.97	
All to chi	Peak	11232.84	4	144068.49	236757.12	
All to chi	Weekend	8696.71	4	115934.92	266750.57	

## Growth Plans

The Create Growth Plan tool lets you generate new plan files containing predictions of future traffic, based on growth forecasts for the current plan. The predicted traffic growth can be generated in one of three ways.

- Per demand grouping using demand grouping traffic growth rates
- Per demand using demand traffic growth rates

- Per interface using interface traffic growth rates

Regardless of the method, the growth is compounded per period increment identified in the Create Growth Plan tool.

A series of plan files is created for each growth period, and each plan contains a report on the prediction process used to generate it. Each newly created plan shows the effects of the traffic growth visually in the plot using default traffic utilization colors, as well as in the Traffic column in the Demands table. By automatically running a simulation analysis on the new plans, you can incorporate worst-case performance into planning decisions.

## Demand Grouping Traffic Growth

Growth plans for network traffic, such as those produced either by trending analyses or by financial or sales forecasts, are typically created for traffic aggregates. You can specify these aggregate growth forecasts or trending estimates as demand grouping traffic growth information. The Create Growth Plans tool allocates traffic to demands to match the aggregate traffic growth specified for each demand grouping. If two demand groupings contain a common demand (for example, a grouping of demands sourced from one site, and a grouping of demands destined for another site), this demand traffic is balanced between the requirement to match each of the demand grouping traffic aggregates.

WAE Design supports two methods of creating growth plans based on demand groupings: one is based on a single period, while the other is for multi-period traffic forecasting.

### Single-Period Demand Grouping Traffic

The basic single-period method grows demands groupings based on the values identified in the Growth Traffic Total, Growth Traffic (%), and Growth Traffic Inc (increment) columns of the Demand Groupings table in the current plan. This approach is useful when all traffic in a given demand grouping is forecast to grow at the same rate or when you want to create growth plans for only one growth period.

If Growth Traffic Total is specified, the base traffic is set to this value. Otherwise, the base traffic is set to the sum of the current demand traffic in the demand grouping. The Create Growth Plans tool then applies the Growth Traffic (%), followed by the Growth Traffic Inc.

Example: If the demands in the demand grouping have a total of 15,000 Mbps of traffic, the Growth Traffic (%) is 5, and the Growth Traffic Inc is 800, the resulting total demand traffic in the first period would be 16,550 Mbps, which is  $(15,000 * 1.05) + 800$ . If you created growth plans for two period increments, the total demand traffic in second increment would be 18,177.50 Mbps, which is  $(16,550 * 1.05) + 800$ .

If the Growth Traffic Total is specified as 20,000 Mbps, then the current demand grouping traffic total is ignored, and the new demand grouping traffic total for the first period would be 21,800 Mbps, which is  $(20,000 * 1.05) + 800$ .

You can populate these growth columns in the Demand Groupings table as follows:

- Choose **Import > Demand Grouping Growth**.
- Choose **Edit > Traffic > Demand Groupings**.
- By editing the Demand Groupings Properties dialog box as described in the following steps.

---

**Step 1** If setting the same growth for multiple demand groupings, select them from the Demand Groupings table.

**Step 2** Double-click a demand grouping, or right-click one and choose **Properties**.

**Step 3** Enter one or more values in the Growth area.

- Traffic Total—Specifies an exact traffic total in Mbps.
- Traffic (%)—Specifies the percent by which to grow current traffic in Mbps.

Example: If the Traffic Total is 30,000 and Traffic (%) is 5, then the growth plan grows the traffic to 31,500 Mbps.

- Traffic Increment—Specifies the amount of traffic to increment by in Mbps.

Example: If Traffic Total is 11,000 and Traffic Increment is 800, the growth plan grows the traffic to 11,800 Mbps.

Growth	
Traffic Total:	<input type="text"/>
Traffic (%):	5
Traffic Increment:	800

**Step 4** Click **OK**.

## Multi-Period Demand Grouping Traffic

The Demand Grouping Traffic table is a way of specifying growth based on demand groupings for one or more specified periods. Example use scenarios could be as follows.

- Traffic sourced and destined for CHI will increase by 10% next quarter
- VPN traffic between MIA and BOS will increase by 1000 Mbps in the third quarter
- Transit traffic will reach 2 Gbps within a year

This table has the flexibility to specify the traffic total, traffic growth in specified increments (for example 10,000 Mbps), or traffic growth at a specified rate (percentage) per demand grouping.

The Demand Grouping Traffic table is created outside of WAE Design (for example, in a spreadsheet). You can then use it, including all of its periods, when creating growth plans. When used with the Create Growth Plans tool, the demand groupings in the current plan must match the names of those you are importing.

For information on creating this external table, see the *Cisco WAE Design Integration and Development Guide*.

## Demand Traffic Growth

This method grows demands based on the value identified in the Growth (%) column of the Demands table in the current plan. This approach is useful when all traffic in the network, all traffic within a given traffic level, or all traffic within a service class is forecast to grow at the same rate. When generating the plans, you have the option of how many plans to create, and each plan will show growth at the same specified rate relative to the previous plan.

Example: If traffic on a demand is 1000 Mbps, if the growth for that demand is defined as 10%, and if you create two growth plans, the first will show 1100 Mbps and the second will show 1210 Mbps on that same demand, representing 10% growth per period.

You can populate the Growth (%) column in the Demands table using either of these two methods.

- Select one or more demands, right-click one, and choose **Properties**. In the Traffic section, enter the growth rate in the Growth % column.
- Import growth rates from a different plan file or table file. For instructions, see the *Cisco WAE Design Integration and Development Guide*.

## Interface Traffic Growth

This method grows measured traffic based on the Growth (%) column in the Interfaces table in the current plan. It differs from demand-based growth methods because measured traffic, rather than demands (simulated traffic), is grown. This is useful for basic forecasting of the current network state, rather than worst-case failure states, which require simulation. When generating the plans, you have the option of how many plans to create, and each plan will show growth at the same specified rate relative to the previous plan.

You can also use this method in conjunction with Demand Deduction to create demands from the forecast measurements.

Example: If traffic on an interface is 525 Mbps, if the growth for that interface is defined as 5%, and you specify three growth plans, the first will show 551.25 Mbps, the second will show 578.81 Mbps, and the last will show 607.75 on the same interface, representing 5% growth per period.

You can populate the Growth (%) column in the Interfaces table using either of these two methods.

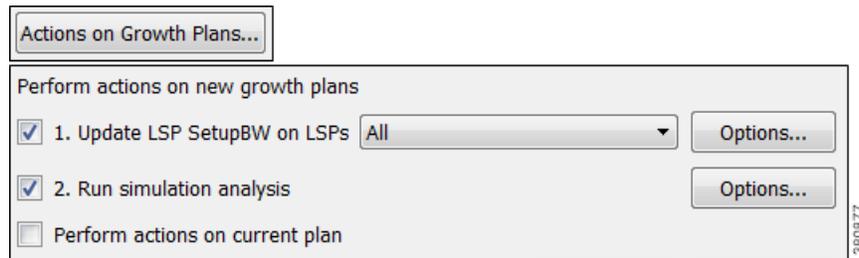
- Choose **Edit > Traffic > Measured Traffic**. In the Growth (%) column of the Interfaces table, enter the percent by which you want to grow the interface traffic for each applicable interface.  
To set growth rates for multiple interfaces to the same value or to random values, select one or more interfaces, right-click, and choose **Set Value**. You then have the option to set fixed values for all selected interfaces or have WAE Design randomly set them within the range that you specify.
- Import growth rates from a different plan file or table file. For instructions, see the *Cisco WAE Design Integration and Development Guide*.

## Creating Growth Plans

- 
- Step 1** Choose **Tools > Create Growth Plans**.
- Step 2** Specify one of the following growth methods:
- Demand Groupings—For information on how to set demand grouping growth based on fixed, incremental, or percents, see [Demand Grouping Traffic Growth](#).
  - Demand Growth (%) in current plan—For information on how to set this percent, see [Demand Traffic Growth](#).
  - Interface Growth (%) in current plan—For information on how to set this percent, see [Interface Traffic Growth](#).
  - Demand Grouping Traffic Table—Browse to or enter the name of the file containing the Demand Grouping Traffic Table. For information on how to create these tables, see the *Cisco WAE Design Integration and Development Guide*.
- Step 3** If creating growth plans based on demand groupings, demands, or interfaces, enter the number of plans to create and the increment by which you want to create them.

- Period increment—Number of percentage increments per period. For example, if a demand is set to grow at 10%, and the period increment is 2, the demand will grow by 21% in the first growth plan.
- Number of periods—Number of growth plan files to create.

Plans are named using the period increment. For example, if you enter 4 as the increment and 2 as period, WAE Design generates two growth plan files: one ending with \_4 and one ending with \_8.

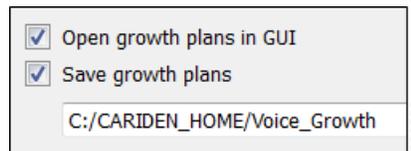


**Step 4** Click **Actions on Growth Plans** to perform these two tasks on the newly created growth plan files. To further specify how to execute them, click their associated **Options** button. You can also choose to execute these tools on the current plan.

- Upgrade LSP setup bandwidth on LSPs—For more information, see [RSVP-TE Simulation](#).
- Run simulation analysis—For more information, see [Simulation Analysis](#).



**Step 5** By default, WAE Design uses the current plan filename as the root name and adds the extensions generated by Step 3 for each growth plan. If you want the new plan files to use a different root name, enter that name in the New plans named field.



**Step 6** To open these new plans upon exiting the dialog box, check **Open growth plans in GUI**.

**Step 7** To automatically save these new plan files, check **Save growth plans**, and then enter or browse to the directory in which you want to save the files.

**Step 8** Click **OK**.

