



Universidad
Politécnica
de Cartagena



PLANIFICACIÓN Y GESTIÓN DE REDES

GRADO EN INGENIERÍA TELEMÁTICA
GRADO EN INGENIERÍA EN SISTEMAS DE TELECOMUNICACIÓN

CURSO 2018-2019

Lab work #2. Network analysis with Net2Plan

(1 session)

Author:

Pablo Pavón Mariño

1 Objectives

The goals of this lab work are:

1. Use Net2Plan to analyze the cost vs. performance trade-offs in an IP network running OSPF protocol. Performances observed will be the per IP flow end-to-end latencies, flow blocking and link congestion situations, under different failures and traffic shifts.
2. Gain practical experience using Net2Plan built-in reports like the availability report, and functionalities like what-if analysis.

2 Duration

This lab work is designed for one session of three hours.

3 Evaluation

This lab work has been designed to guide the students in their learning of the subject. The annotations the students make in this document are for their use when studying the course, and do not have to be delivered to the teacher for evaluation.

4 Documentation

The resources needed for this lab work are:

- Lecture notes of the course.
- Net2Plan tool and their documentation (see <http://www.net2plan.com/>).
- Instructions in this wording.

5 Previous work before coming to the lab

- Read lecture notes related to risk & failure analysis.
- Refresh your reading in <http://www.net2plan.com/>, in particular, on the availability report.

6 What-if analysis

The **what-if** analysis functionalities of Net2Plan, permit the user observing the immediate effects of failures and traffic shifts in the network.

What-if analysis is based on a built-in or user-developed *network reaction algorithm* (in Net2Plan, also called *provisioning algorithm*), that codes how the network will react to these changes. In this lab session, we will use a Net2Plan built-in algorithm that implements the reaction of IP/OSPF protocol.

Follow the next steps:

1. In the **What-if analysis** tab, read the description in the screen about the main guidelines of what-if analysis. Activate the **what-if** functionality by clicking the toggle button **Toggle What-if Mode**. This enables the selection of the reaction algorithm in the panel below.
2. In the **BuiltInExamples.jar** file, select the reaction algorithm **Online_evProv_ip0spf**.

From now on, everytime a demand offered traffic is changed, or a link/node changes its state from up to down or viceversa, the reaction algorithm is called and the traffic is rerouted according to OSPF.

6.1 Check the reaction to failures

1. Open **Net2Plan** and load the network design in the file **example7nodes_withTraffic.n2p**.
2. In the **View/Edit network state** tab, **Network** subtab, change the layer name to **IP** in the **Layer information** table.
3. Observe the current traffic routing in the non-failure state. Is all the traffic carried? Is there any oversubscribed link?

4. Set as down the two directions of the bidirectional link between Valencia and Madrid. Is all the traffic carried? Is there any oversubscribed link? Did the worst case latency of the flow from Sevilla to Barcelona improved? why?¹

5. Set as down the two directions of the bidirectional link between Murcia and Malaga. Is all the traffic carried? Is there any oversubscribed link?

6. Set as down the two directions of the bidirectional link between Barcelona and Zaragoza. Which fraction of the traffic is not carried? Is there any oversubscribed link?

¹Look in the topology panel at the paths followed by the flow in both situations, by clicking it. Notice the ECMP effect in the non-failure state.

7. Set all the links as up. Then, set the node in Madrid as down. Which are the affected demands? Is there any affected demand, aside the one initiated or terminated in Madrid? Which is the path of the traffic from Zaragoza to Sevilla now?

Quiz 1. The student is encouraged to observe the reaction to failures under arbitrary combinations of links and node failures.

6.2 Check the reaction to changes in the links' capacities and traffic shifts

Set the network in the non-failure state. Change all the link capacities to 15, by selecting all the links in the **Links** table, and using the right-click option **Set selected links capacity**.

1. Are there oversubscribed links? Are there any demand traversing oversubscribed links?

2. Is OSPF making any change in the routing to deviate the traffic from oversubscription?

3. Set again link capacities to 50 Gbps. Now, manually set the traffic from Barcelona to Madrid to 80 Gbps. Is there any link oversubscription? Is OSPF making any change in the routing to deviate the traffic from oversubscription?

6.3 Check the reaction to changes in the IGP weights

The OSPF weight of each IP link can be set manually by setting the attribute `igp-linkWeight` with the desired IGP cost (greater or equal to one).

1. Load the network design in the file `example7nodes_withTraffic.n2p`.

2. In the **Links** table, select all the links, and use the right-click option **Attributes -> Add attributes to selected elements**, to add the attribute `igp-linkWeight` with value 10.0 to all the links.
3. Set the capacity of all the links to 20. Note that the IP link between Valencia and Madrid has a 99% utilization.
4. Force the update of the OSPF reaction algorithm, by setting up and then down any link. This is needed since Net2Plan does not automatically call the reaction algorithm when attributes are modified.
5. Use the option **Attributes -> Expand all attributes**, to have each defined attribute in a different column.
6. We are interested in reducing the utilization of the Valencia-Madrid links. For this, change the IGP weights manually of the links between Valencia - Madrid, to 11. Is the utilization of those links reduced? Did that produce a solution without oversubscribed links?

Quiz 2. Set the capacity of all the links to 50. Now, make some manual tests changing IGP weights to find a solution with the lowest congestion in the non-failure state. Note how difficult is to find good IGP weights in this manual form. Recall that this difficulty is even higher if we want to optimize the congestion not only in the non-failure state, but also considering multiple failure states.

7 Failure analysis

This section is devoted to make use of the Net2Plan built-in availability report, to analyze the network health under different failures.

Load the network design in the file `example7nodes_withTraffic.n2p`. Now let us assume that we are willing to design an IP network with the following characteristics:

- The SLA associated to each traffic demand requires for each flow to not traverse IP links with an utilization above **90%**, and to have a maximum end-to-end propagation latency of **7 ms**.
- The network should be fault tolerant to all single IP adjacency failures. That is, maintain the SLAs of all the flows if the IP links in both directions between two nodes, simultaneously fail.
- The network should be as robust as possible for multiple IP adjacency failures.

This section will use Net2Plan to analyze the health of the network in this respect, and assist the search of valid network designs. For this, we will make extensive use of the built-in Net2Plan report called `Report_availability`. This report, automatically generates multiple failures in the network, applies a *reaction algorithm* (or *provisioning algorithm*) to simulate the network reaction to each failure state, and compiles statistics for each demand and for each link, for each failure state tested.

7.1 Defining the SRGs, and assessing the single-SRG failure health

The *Shared Risk Groups* (SRGs) are the elements defined in the network to represent the *risks* of a user-defined set of failures to happen. Each SRG has associated:

- The set of links and nodes that simultaneously fail of such risk happen.
- The statistical probabilities associated to the risk: the Mean Time To Fail (MTTF, the average time between a the failure is repaired, and happen again), and the Mean Time To Repair (MTTR, the average time needed to repair this failure when it happens.)

Follow the next steps:

1. In the **Demands** table, set the maximum acceptable end-to-end latency for each demand to 7 ms, using the right-click option **Set maximum e2e limit to selected demands**. Is any demand outside this limit in the non-failure state?

2. In the **SRGs** table, use the option **Add SRG from model**, to add one SRG for each potential IP adjacency failure. For this, use the **PER_BIDIRECTIONAL_LINK_BUNDLE** option. Set an MTTF of 8760 hours and an MTTR of 12 hours. Check out that 8 SRGs are created. Which is the availability of each IP adjacency according to this model?²

3. Load the report **Report_availability** in the **View reports** tab. Read the description, and what each of the parameters is for.
4. Load the **Online_evProv_ipOspf** algorithm as the provisioning algorithm (**provisioningAlgorithm_file**). Use the report for analyzing the network fault tolerance for single-SRG failure (**analyzeDoubleFailures = false**), for those SRGs already defined in the network (**failureModel = SRGfromNetPlan**).

According to the analysis report:

1. Is there any failure state where there is traffic blocked?

²To observe this number, increase the number of precision digits shown, making right-click in the column header, and then **Double format -> 5**.

2. Is there any failure state where any IP link is over 90% utilization?

3. Is there any failure state where the end-to-end latency of any flow is higher than 7 ms? If so, which are the affected IP flows, and which failures make them violate the SLA?

7.2 Changing the topology to fit latency requirements

To address the SLA violation respect to the end-to-end latency requisite, a new IP adjacency will be added between Madrid and Barcelona. For this:

- Add in the topology panel a bidirectional link between those nodes, pressing **SHIFT** key, while dragging the mouse from one node to the other.
- Then, in the **Links** table, set the link capacity to 50, and set the appropriate link length using the right-click option **Set selected links length -> As the haversine node pair distance**. This will set the length to 504.65 km.
- Remake the SRGs, one per IP adjacency, to include the new IP adjacency.
- Save the new design with the name **plusAdjacency.n2p**.

Remake the report analysis.

1. Are all the SLAs met now, in all the single IP-adjacency failure states?

2. Observe in the **General** sub-tab in the IP layer tab, the information of the average end-to-end number of hops and average end-to-end latency, before and adding the IP adjacency. Using this information, and the analysis report, fill in the following table.

Effect of adding Mad-Bcn IP adjacency

Concept	Original design	After adding IP adjacency
Worst case latency any failure		
Worst case oversubscription any failure		
Worst case blocking any failure		
No failure state av- # hops		
No failure average e2e latency		
No failure worst case e2e latency		

- Why the Malaga-Barcelona flow has a worse end-to-end latency now? Note that OSPF splits the traffic into the shortest paths, and the worst case latency value corresponds to the worst among the shortest paths.

Quiz 3. Try adding other adjacencies, to address the end-to-end latency problem, and analyze their performances.

7.3 Analyzing the robustness for double failures

We are interested in this section in analyzing the robustness of the network under single and also double IP adjacency failures. In our case, this corresponds to single and double SRG failure.

Follow the next steps:

- Load the file `plusAdjacency.n2p`.
- Fill in the following table, running the `Report_availability` report twice. One with the same options as before, and other checking now the `analyzeDoubleFailures` option. In the latter case, the report will simulate the network reaction in (i) the non-failure state, (ii) all the single IP adjacency failure states, and (iii) all the double IP adjacency failure states.

Failure analysis

Concept	Single-failure	Single & double failure
Prob. of not-considered failures		
Worst availability among all IP flows		
Worst survivability among all IP flows		

Answer the following questions:

- If we are interested in guaranteeing a survivability of 5-nines (99.999%), can the single-failure analysis be enough? why?

- According to the more accurate availability estimations including single and double failures, is the network able of providing a 5-nines availability?

3. Which are the only flows surviving to all double failures? are all the in and out nodes of those IP flows of at least IP degree 3 (i.e. connected to at least other 3 nodes)? why?

Quiz 4. Compute the network robustness figures, in the single-SRG case, for the same case as before but adding also node failures occurring one every two years, and with a repair time of 5 minutes.

8 Capacity & cost analysis

In this section, we are interested in analyzing the capacity requirements in the network, comparing:

- **Non-fault tolerant design.** The IP capacity needed if fault-tolerance was not a requisite. This corresponds to the capacity needed to meet the SLAs in the non-failure state.
- **Fault-tolerant design.** The IP capacity needed, if single-SRG fault-tolerance was required, but there is no requisite to have 5-nines availability. To obtain these values, run the `Report_availability` focusing just in single-SRG failures. Then, in the per-link table, we observe the worst case traffic that is being traversed by each IP adjacency, in the worst simulated failure state.
- **Fault-tolerant & robust design.** The IP capacity needed, is single-SRG fault tolerance is required, AND 5-nines availability is requested for all IP flows. To obtain these values, run the `Report_availability` as before, but checking also the `analyzeDoubleFailures` checkbox. Now, the per-link table shows the worst case traffic that is being traversed by each IP adjacency, among the non-failure, single failure and double failure states.

Fill in the following table, with minimum cost required capacities considering that:

- Use the design saved as `plusAdjacency.n2p`.
- All links should have an utilization below or equal than 90%.
- All the link capacities must be composed of modules of 10 Gbps (cost 1), or modules of 100 Gbps (cost 3.5), but not modules of both in the same adjacency.
- Since the traffic is symmetric in this example, count the number of 10G/100G modules independently, and sum them³.

Capacity & cost analysis

Concept	No requisite	Fault tolerant	FT & robust
# 10G modules			
# 100G modules			
Total cost			

Quiz 5. Repeat this cost analysis considering a design where other adjacency was added to the baseline design, instead of that between Madrid and Barcelona, for having a design meeting the end-to-end latency constraints.

³When the traffic is assymmetric, the capacity to allocate in an IP adjacency is computed from the largest between the traffic in both directions.

Quiz 6. Repeat this cost analysis considering a design where the latency constraint is set to 10 ms, and thus there is no need to add an IP adjacency to the original design. Comparing with the results in this section, which are the eventual cost savings of relaxing the latency constraints?

9 Work at home after the lab work

The student is encouraged to complete all the tasks that he/she could not finish during the lab session, and the Quizzes proposed.