

Rapid Spanning Tree Simulation Guide 0.06d

Roger Pfister

Mick Seaman

The Rapid Spanning Tree Protocol simulation first distributed at the November 2000 IEEE 802.1 meeting functions as a Visio™ 'Solution'. Just as Visio allows you to draw Organization Charts using a special 'Stencil' of shapes, the Rapid Spanning Tree solution allows you to draw and simulate switched networks to study the behavior of RSTP. It shows the transmission of BPDUs, Port Roles, Port States, and internal state machine variables associated with each Bridge Port. It allows links (LAN Segments) to be broken and restored, and bridges to be powered on and off, to help you validate the expected behavior of the RSTP state machines, rapidly reconfiguring the active topology to restore connectivity without introducing temporary loops. Scenarios can be calculated, recorded, and replayed, navigating both forward and backward in time to facilitate study of the protocol.

This guide describes version 0.06d which is a simulator document P802.1wD8.

This guide tells you how to acquire, install, and run the latest simulation tool distribution.

Purpose

This simulation tool is being made available 'as is' to facilitate study of IEEE P802.1w Rapid Reconfiguration and RSTP and advance the standards process. There is no commitment on the part of the authors to further develop this tool, to fix real or perceived defects, or to continue to distribute the tool in the future. This statement should not be taken to represent any disavowal of or reduction in the intellectual property rights of the authors or their employers.

Unsupported Software

Except as otherwise explicitly and personally communicated, the authors have no time or resource to answer questions about the tool itself, accept bug reports, deal with enquiries as to if there is a new release or further planned development, or respond to any other input on how the program could be improved. You may receive no response to any such communication.

Observations and constructed scenarios to facilitate discussion of the behavior of RSTP, the faithfulness of the tool in reproducing such behavior, and to advance the standards effort consistent with the 'Purpose' described above can be emailed to Mick Seaman mick@telseon.com or circulated to the P802.1w working group directly.

Permitted Use

Permission is granted to use the simulation tool for the 'Purpose' described above. In using the tool you agree (a) not to reverse engineer or copy the design of the tool or its methods without the express written consent of its authors (b) that you are aware of the restrictions and possible shortcomings of the tool described in 'Disclaimer' below (c) that the 'Activation Icon' provided in the tool is retained and displayed in all diagrams created using the tool and subsequently distributed.

Disclaimer

While the authors are interested in advancing standards in general and IEEE P802.1w Rapid Reconfiguration in particular, there can be no guarantee that the Rapid Spanning Tree simulation represents an accurate or permitted interpretation of any revision of P802.1w or the final standard when that is approved. No representation is made as to the fitness of the tool for any particular purpose, including but not limited to the design and performance modeling of networks. The tool is being distributed (see Purpose above) to facilitate study of RSTP by participants in the standards process. It is the responsibility of the users, not the authors, of the tool to verify the correspondence between the behavior represented by the tool and the state machines and other provisions of P802.1w.

Copyright

Submissions (slides, printed papers, electronic versions of these) to the IEEE standards process are copyright IEEE. The templates, i.e. the initial drawings that you will see when first selecting the Rapid Spanning Tree solution in Visio, are no exception as will be any drawings that you produce and share with the standards group. This copyright may extend to the literal representation of the application as a collection of bytes, allowing the IEEE and participants in the IEEE standards process to make copies as required for study. However these rights do not extend to use, modification, reverse engineering, or reuse of the simulation tool or its internal methods, anymore than they extend to the coding and representational methods used in the postscript, PDF, word, or Zip files used to share text and diagrams. Use of the simulation tool is subject to the conditions described in 'Permitted Use' above; with these restrictions the copyright of any diagrams and scenarios that you create using the tool rests with yourself in the first instance.

Requirements

To run this software you will need any version of Visio 2000 (US or International English) running under Windows 2000 (preferred), Windows NT4.0, or Windows98. As this software was built and tested on Visio™ 2000 (SR1), SR1 is advised. It is your responsibility, not ours, to ensure that you are running a legally licensed copy of Visio. We do not distribute Visio.

Distribution

Releases of the simulation tool for standards contribution and study are distributed by posting a .zip file in the P802.1 documents directory. Go to:

<http://grouper.ieee.org/groups/802/1/index.html>

and select the documents link and the directory for the year of distribution (currently docs2000). Each distribution file has a name of the form RVR_NNan.zip where 'NN' is a sub-version number, 'a' is a single alpha character and 'n' is an optional revision number. At the time of writing the current release is 0.06c and is distributed in RVR_06c.zip.

Unzipping the release file should extract the following or similar files:

- Rapid Spanning Tree.vsl
- Rapid Spanning Tree.vst
- Rapid Spanning Tree.vss
- Rapid Spanning Tree Rarely Used.vss
- Rapid Spanning Tree Simulation Guide.pdf
- Read_me.txt
- Permitted_use.txt
- install.bat

The .vsl file is a Microsoft DLL and contains the bulk of the intelligence of the tool (see below). The .vst is a Visio 'Template', i.e. the drawings that come up when you open a 'New' file. The .vss is a 'Stencil'. This is the set of prepared shapes, connectors, and icons that you can drag and drop onto the networks you draw and are recognized by the simulation tool.

The 'rarely used' stencil contains all additional objects that can be used to control the simulation. The standard template loads and show both of the above templates. The distribution may also contain a number of test scenario drawings as .vsd files. The Read_me.txt and Permitted_use.txt files are also included by convention. They reiterate information found elsewhere in this guide.

Installation

Prior to installation ensure that Visio is not running.

Installation is simply a matter of copying files to certain directories used by Visio. The .vsl file should be copied to the 'Startup'¹ directory and the .vst and .vss files to a 'Rapid Spanning Tree' sub-directory in the 'Solutions' directory. When first installing the simulation tool it will be necessary to create the 'Rapid Spanning Tree' sub-directory with exactly that name.

The accompanying read_me.txt, permitted_use.txt, and install.bat files can safely be kept in the Rapid Spanning Tree directory, though of course the purpose of install.bat is to copy the files, creating the directory if necessary.

The example scenarios '.vsd files' (if any) can be copied to any convenient directory. They will automatically pickup the necessary stencil from the Rapid Spanning Tree directory.

The install.bat file included in the distribution can be used to automate installation of the Rapid Spanning Tree solution. It assumes that Visio is installed on the current drive, though it takes a single optional parameter x: to copy to drive x instead.

Versioning

The stencil provides objects to be placed on each drawing page that display the Rapid Spanning Tree .vsl version, the stencil (.vss) version, and the version of the template (.vst) file used to create a drawing (.vsd) file.

Once you have opened a Rapid Spanning Tree drawing you should check, by eye, that the stencil version matches the .vsl version e.g. 0.06d, 0.06d. Differences here means that you have not correctly installed or upgraded from a previous version

Uninstalling Rapid Spanning Tree

The Rapid Spanning Tree solution can be completely removed from Visio by removing the .vsl file from the Startup folder while Visio is not running. It is not necessary to delete the Rapid Spanning Tree subfolder from the Solutions directory, although that can of course be done.

¹ Normally Visio locates its 'Startup' and 'Solutions' directories relative to the location of the Visio32.exe file itself. In the default install they will be found at C:\Program Files\Visio\Startup and C:\Program Files\Visio\Solutions. If you have a non-standard install consult Tools>Options>File Paths from within Visio to find which paths are being searched.

Starting Visio

Once the above files are in place you can run Visio and select "Rapid Spanning Tree" from the list of possible drawings. This will invoke the Rapid Spanning Tree template (.vst file) and display the first of a number of example drawings.

Alternatively you can simply open (click or double-click) one of the example .vsd files.

In both cases you should now see a new toolbar with 15 or so buttons, including a "Red Diamond" roughly in their center. Their use is explained below. If you do not see the Toolbar you will not be able to use the simulation. There are two common explanations for a missing toolbar. The "Rapid Spanning Tree.vsl" file may not been copied to the Startup folder (see above), or the toolbar may have been closed by clicking its corner 'x' while it was floating. In the latter case it may be restored. Point to the gray area where toolbars normally live. Right click the mouse to reveal a menu whose last item is "customize" and click that. You should now see a dialog with Rapid Spanning Tree as the last entry. Tick that entry. If it does not exist then installation has not been completed successfully². Exit and restart Visio, and if the problem persists check the installation files³, directories, and paths as described above³.

Getting Started

Once you have opened the Rapid Spanning Tree template or one of the supplied .vsd files and can see the Rapid Spanning Tree toolbar, click on the 'Red Diamond' in the middle of the toolbar. This will recalculate the scenario (a single page drawing of bridges and their connectivity together with the changes made at specific times).

Once the calculation is complete all the buttons on the toolbar will be active, and pointing to them will give you an idea of what they do. Experiment with the leftmost 'Play' buttons. You will see the simulation time and the 'step number' change as you play through the calculated simulation. Zoom in and out or scroll up and down to familiarize yourself with all the objects on the page. Experiment with moving the red 'Port Probe' object from port to port to see the current state machine variable values.

Try adding new bridge and hublet objects from the template and connecting them with LAN segments. You will need to press the 'Red Diamond' again to recalculate the simulation, before playing through it. Once you have created a new scenario, the Visio drawing can be saved as a .vsd file. When the file is reopened the calculated steps will have been lost, but can be recalculated exactly as they originally occurred.

² If you have created the page from scratch check that all the mandatory items are present (see below).

³ Visio can get into a very confused state following non-standard installs, and there may be nothing for it but to uninstall and reinstall Visio, borrow someone else's machine, or as I did (for reasons unrelated to RSTP) move to a clean Windows 2000 machine. On NT Visio does not appear to uninstall cleanly, but leaves traces in the registry which appear to impact reinstalls.

How it works

A basic understanding of how the Rapid Spanning Tree simulation works is helpful to even the casual user.

The simulation proper is independent of Visio⁴ and does not rely on Visio's representation of bridges, ports, LAN segments (links), or hubs as it runs. At the beginning of time the Visio diagram is scanned, auto-magic assignments are made to configuration parameters not set by the user, and the simulation's own bridges, ports, and connectivity instantiated⁵.

The simulation is then used to 'calculate' a number of steps, each step being defined by the occurrence of a significant event – a change in port state, a BPDU transmission, a change in a significant variable. Some of the parameters that define a step may be controlled by the simulation user. At each step the simulation's state including the RSTP state machine variables and the network connectivity are recorded.

Once a step has been calculated it can be displayed or 'played'. The simulation's own variables are reflected back into the Visio drawing to show the state at each step.

Because the simulation state is available at each step, the user can calculate forward from any given step without having to recalculate from the beginning. The simulation has saved its own variables at each step, not a copy of the Visio diagram at that step. If the diagram is rearranged without changing the network elements or their connectivity, the changes will apply to all the steps. The simulation proper is unaware of the actual layout of the diagram.

The link between the Visio diagram and the simulation proper has to be told about changes in connectivity. This is done by pressing a 'change recorder' button on the toolbar, making the change, and then pressing the button again. A visible record of all changes and their timing is kept. Although the simulation proper can accommodate all sorts of changes while running, the creation of a new bridge or bridge port for example, the change recorder cannot currently accommodate all events. If one occurs that it cannot handle (see below for a list of its current capabilities and exceptions) the calculated steps are discarded and the scenario must be recalculated from the beginning.

⁴ Those familiar with Visio and Microsoft technology seem to expect the simulation to be written in Visual Basic, of which there is almost none. The little there is simply checks that the .vsl add-on has been loaded. The simulation proper is written in C++, and is a version of the C based simulation used for the development of LACP (see IEEE 802.3ad Link Aggregation). This in turn was based on the GARP and GVRP development environment. The translation between the state machines in P802.1w and the simulation code was largely achieved by cutting and pasting the C++ state action statements and transition conditions into a trivial structure that represents the interconnection of the states. The purpose of this exercise is to verify that the description in the standard provides useful protocol behavior with the minimum of reliance on implementation expertise, rather than to display the authors' code crafting skills.

⁵ I do not claim that this is necessarily a better way of proceeding than the 'obvious' one. It proved to be very convenient in this case, and has a significant impact on the way the simulation is best used, hence the current description.

Creating a new scenario

You can drag, drop, and connect all the necessary objects from the full stencil onto a new page. However it is easier to create a new drawing using a subset stencil as follows. Open one of the distributed test case .vsd files and use Save As to rename a copy with your chosen file name. Insert a new page, visit an old page, select and copy all (Ctrl-A followed by Ctrl-C), return to the new page and paste (Ctrl-V), and finally line up the copied material on the page (paste does not put it down exactly where it is wanted). You can then rearrange, add and delete bridges and links as necessary for the new scenario before saving your file once more.

If you are creating a new page from scratch or rearranging objects on a page to suit your own layout preference you need to ensure that the following objects are present:

- 'Activation Icon'. This draws attention to the permitted use and limitations of the simulation tool and must be in from the edge of the page, in front of any overlapping objects, and unchanged.
- ".vsl Version"
- "Time" and "Step x of n" displays
- "Activity List" and the "Change List".

Without these objects the simulator will not operate and the 'Red Diamond' toolbar button will remain gray. After all the necessary objects have been added switch to any other page and return to the new page to prompt the simulator to start.

Many other optional objects can be placed on the page. Some provide input to control the behavior or the output of the simulation, some display output, and some do both. Almost all controlling objects are only read when one of the toolbar recalculation buttons, the 'Red Diamond' for example, is clicked.

The Toolbar

The Toolbar buttons allow you to:

- Rewind
- Calculate
- Play

a scenario.

When a scenario is rewound any changes prior to the current time are 'unplayed', the current time is set to zero, the Visio diagram and its control icons are scanned to auto-magically assign any missing network parameters, and the simulation itself is initialized.

A scenario may be calculated:

- Step by step
- To a specified time
- To 'stability' (see below)

The time specified may be set to be later than the tools own interpretation of stability. This is useful if a change is to be recorded after the topology has stabilized.

Individual buttons combine the three basic elements of simulation navigation (reset, calculation, play) for ease of use. Currently the buttons can be considered in three groups.

The left hand five buttons (dis)Play the calculated steps in the scenario:

- Bell first step
- Left Arrow back one step
- Hourglass step nearest the time box time (set first)
- Right Arrow forward one step
- Eight Ball last calculated step

The next three buttons **Rewind, recalculate to stability⁶ and display:**

- Red Diamond the first step
- Hourglass the step nearest the time box time (set first)
- Eight Ball the last step (stability)

The next five buttons **Calculate forward⁷ and display:**

- Bell first step
- Left Arrow back one step
- Hourglass the exact time in time box (set first⁸)
- Right Arrow forward one step⁹
- Eight Ball to stability

⁶ Any steps previously calculated after the step determined by the simulation to be 'stability' are discarded.

⁷ The first two buttons are just copies of the Play functionality, since no forward calculation is required.

⁸ Will calculate a step at the target time, even if that step would not normally exist.

⁹ Will wipe out any subsequent steps previously calculated.

Stability

A scenario is said to have reached stability (as defined by the calculation logic) when there has been no change of Port Role or Port State for the "Stability Secs" selected. This defaults to 16 seconds (Forward Delay plus one second). Once stability has been detected by the recalculation engine all but the first 3 seconds of activity following is discarded, there being little benefit to viewing purely repetitive behavior.

Making and recording changes

First navigate to the desired time¹⁰ for the change. If this is not at a step go to an earlier time, change the displayed time to the desired time and use the 'calculate – to a TIME' button. Choosing a specific time to make the change is a good idea if the scenario is to be run with different random seeds to explore the effects of variable process and transmission scheduling.

Before editing the diagram to make the change(s), click on the 'connection change recorder – On/Off' toolbar button. The following changes can now be made:

- moving the ends of LAN segments¹¹
- powering on or off bridges

Any other changes reset the scenario and require complete recalculation. Any components required at any time need to be on the initial diagram¹².

As you record changes you will see them being stored in the 'Change Display' box. If you start recording at the same time as existing changes your new changes will be added to the end of the list. Any changes set 'at a future time' from the time of a recording will be discarded. This can be used to selectively wipe out changes. A separate toolbar button allows you to discard all recorded changes.

Once all the changes at a given time are complete click on the 'connection change recorder – On/Off' toolbar button once more. Toolbar buttons that were grayed while the change recorder was in operation become active once more. Use the calculate buttons on the toolbar to calculate forward once more.

As you play the scenario the changes you have recorded will also be played or unplayed (if you are stepping backwards) and the consequent topology changes should become apparent.

¹⁰ Changes cannot be made at time zero.

¹¹ It is a good idea to use spare hublets to terminate LAN segments that are 'broken', i.e. do not connect two bridges. While a LAN segment can be left at an arbitrary position in the page, that position rather than a connection to a movable object has to be recorded which limits your flexibility to rearrange the diagram.

¹² Bridges can be left powered off until required, hubs and LAN segments are purely passive so can easily be left around until needed.

Powering on and off

A bridge can be powered off and on by selecting the bridge and right-clicking to bring up a menu that will allow you to toggle the power. Power toggling events are recordable (do not forget to turn recording on) and so can form part of a scenario that can be easily replayed. A powered off bridge is shown in dotted outline. Its ports are still connected but are held in Disabled state.

Setting a non-zero 'Rand Seed' value causes each Bridge to be powered on at a random (but repeatable) time within one HelloTime of the simulation start i.e. within the first two seconds. Different seeds will generate starting sequences. Use of this feature ensures that bridges do not continually operate in lock step timing and provides for more accurate simulation of protocol behavior. However it causes many more simulation steps to be generated for scenarios with large numbers of bridges and thus greatly increases calculation time. It is best to run with a random seed of zero while constructing your intended network scenario, and then to switch to non-zero values for simulation runs.

To avoid confusion between the two power control mechanisms you should not manually toggle the power within the first two seconds if random power up is to be used.

Selecting forceVersion

The simulator allows you to set the forceVersion parameter for a Bridge. Select the Bridge and right click to bring up a menu that will allow you to toggle forceVersion. Such changes are not recordable and must be followed by a reset (rewind and recalculation), so forceVersion is a constant throughout the life of a bridge. A bridge with forceVersion 0 is displayed with a blue outline.

Managing Spanning Tree Parameters

To help you construct scenarios quickly, the unique identifiers required by the spanning tree algorithm are 'auto-magically' assigned. A non-printing hexagonal 'Walk From' object can be used to guide the assignment algorithm by attaching it to any bridge port using a LAN segment connector. This helps you to simulate a number of different bridge identifier assignments in a given network without having to manually configure all the bridges. If there are separate networks on a scenario page several 'Walk From' objects can be used, the number in the center of each controls the assignment order.

The simulation tool does allow you to explicitly manage individual Bridge and Port Identifiers, and to override the default settings for Port Path Cost and Port Priority¹³. First select the box that surrounds the Bridge or Port Identifier (as appropriate). Once selected it will appear surrounded by tiny padlocks. Right click the mouse and select Format>Special from the cascading menus. This will bring up a dialog box with three settable data items.

The Bridge Identifier is set in Data1 using up to three hex characters. The Port Identifier, port Path Cost, and Port Priority are set in Data1 (two hex characters), Data2 (a decimal number), and Data3 (two hex characters) respectively. Clearing the fields reinstates the default or the auto-magic assignment and the latter will resolve any conflicts it finds¹⁴.

The scenario has to be rewound and recalculated before any changes take effect. The simulation tool does not currently support on the fly modification of spanning tree management parameters.

OperEdge and OperPointToPointMac

Rapid reconfiguration is facilitated by the identification of ports at the edge of the switched network that are not connected to other bridges, and by handshakes¹⁵ between bridge ports connected by point-to-point links.

The simulation sets operEdge true for a port if and only if it starts at time zero with a LAN segment attached to it that LAN is not attached to any other object or is only attached to WalkFrom object.

The simulation sets operPointToPointMAC true if operEdge is false and the port is currently connected to a LAN which has no more than one other port attached.

Port Probe

Once a recalculation button has been pressed you can point the red Port Probe into a port and view the state machine variables in the Probe Output box. The Port Probe can be moved from port to port as needed without further recalculation.

¹³ The defaults are also manageable. Use the stencil object which you find already on or below the page in the supplied template and test case files.

¹⁴ Sometimes in a surprising way.

¹⁵ Proposals inviting Agreements.

File Output

When a recalculation button is clicked a results file is produced. You can control the path and the name of the file. The file name is taken from the 'Scenario Name' control object with '.txt' appended.

Step Execution

In theory the operation of each RSTP state machine is arbitrarily quick and arbitrarily interleaved with the operation of all the other state machines¹⁶. To gain insight into what is actually happening it is however useful to define and know the execution order actually used, and to pause to display the results at each 'step'. An arbitrary step to step 'process scheduling' or 'process execution' delay of 1 millisecond is imposed to spread these results over time, simulating a fairly fast distributed bridge implementation.

Here is what happens when the 'Calculate - Forward one step' button is clicked:

1. The simulator checks to see if a new step is required at this point. Has the port state or selected role of any port changed since the last step or was a connectivity change recorded at the last step?
2. Any connectivity changes drawn at the last step are incorporated into the simulation network. If a new LAN has been created, which happens whenever a point-to-point link is broken in two, all LANs will be renumbered.
3. The time is incremented by 1 millisecond.
4. Any 'in transit' BPDUs are delivered, thus simulating a 'one tick' transmit time.
5. The simulation cycles through one pass of its time chain and all the port state machines execute once.
6. Any connectivity changes recorded at the current time are displayed. These changes will not affect the simulation's network connectivity until the next simulation tick (see above), so other changes can now be recorded without invalidating the existing display.
7. The activity list box is updated with the:
 - a. Time and step number and the port states at the end of the step
 - b. List of all the received BPDUs (these may have caused the above states)
 - c. List of all the transmitted BPDUs (these will be received at next tick)
 - d. List of all connectivity changes recorded at the end of the step.

¹⁶ With the exception that the actions taken inside an individual state box are atomic with respect to the operation of any other machine/