# Securing RSTP

Mick Seaman

This note is an preliminary look at using MACsec to secure RSTP. This revision takes advantage of the controls suggested in P802.1AE/D2 to simplify an earlier proposal.

A deployment plan and security policies for use with RSTP/MSTP in conjunction with MACsec are proposed.

## Overview

This note

- Summarizes the challenges and ingredients required for a successful deployment plan.
- Puts the argument for ensuring that standard management encompasses deployment.
- Lists challenges that might occur in a MACsec deployment.
- Suggests a step by step deployment plan that allows these challenges to be encountered and investigated while maintaining network connectivity.
- Describes RSTP policies that ensure that only authorized systems can perturb the spanning tree configuration and learnt address information.
- Considers early deployment of software based MACsec protection for configuration protocols in advance of full rate hardware, with special reference to RSTP.

## Protocol deployment

Network administrators rarely have the freedom to take entire networks of significant size out of service for an extended period to perform simultaneous upgrades of systems. At the same time the number of things that could go wrong on a big bang transition from undeployed to fully deployed is large. So large in fact that a "just try it and fallback if it all doesn't work" approach is both unlikely to succeed and unlikely to provide enough diagnostic data for success on a second or subsequent attempts.

A staged deployment plan is required. Each stage should:

1. be a small step toward the end goal
2. be easily reversible
3. provide positive feedback, and diagnostic data
   a) confirming that the step occurred
   b) pinpointing the cause of any network problem introduced
   c) identifying problems that should be fixed before taking the next step.

In an ideal world each stage can be made risk free, i.e. the network will continue to work as expected, provided any problems identified in the prior stage were addressed. In most deployments the diagnostic capabilities have to be more powerful than would be required if a methodical staged plan had been used. It is rare, for example, that an adequate inventory of network devices and their current state is produced until after a deployment problem has occurred.

## Specifying management

Management controls are clearly needed to support staged deployment. Counters of both normal and error events are required to confirm that each stage is proceeding to plan, and to diagnose failures.

There are alternative ways of introducing MACsec. If these are not explicitly discussed it is likely that the management controls provided by the standard will be selection of those required by a number of these, but insufficient to support any one of them. Whether proprietary management from any vendor will be sufficient is anyone's guess, but the chance of multi-vendor staged deployment is low, so practical interoperability will be confined to those cases where the network has magically sprung into being in its final complete and debugged form.

## MACsec deployment challenges

MACsec secures a network one LAN at a time, which greatly helps deployment[1] by reducing number of inter-system dependencies, particularly if the LAN is truly point-to-point. Deployment might still be disrupted because

1. Stations[2] that provide or require connectivity and are essential to service delivery are

[1] LANs most open to attack can be secured first, and a secure perimeter implemented.

[2] Most likely bridges, routers, or servers.

attached to the LAN but have been missed[3] in the upgrade process.

2. Cipher suite selection has failed.
3. Key agreement protocols fail[4].
4. Authorization levels are incorrectly set or not correctly bound to authenticities.
5. Authentication credentials are not properly distributed or maintained[5].
6. Incorrect client policies have been implemented.

The second of this reasons is largely dealt with by requiring a mandatory default cipher suite for all conformant implementations.

## Full MACsec deployment

The plan to secure RSTP operation on a LAN is somewhat simpler if it can be assumed that full MACsec capabilities are available in each of the systems attached to the LAN, so that will be described first.

We will also assume that the risk of missing stations from the pre-upgrade network baseline process is low. This assumption is reasonable if point-to-point links are being upgraded, but to help it we plan to upgrade one LAN at a time, so when the network breaks we have a good chance of identifying where (provided that not too much else, unconnected to MACsec, is being changed at the same time).

The 'standard' interface stack configuration is used to support this deployment, i.e. both the MAC Relay Entity and the RSTP Entity (in Bridges) are attached to the Controlled Port of the SecY providing the interface to the LAN.

The upgrade steps are:

1. Upgrade all the system attached to the target LAN so they are MACsec capable. The management parameters of the SecY[6] are set as follows:

   Secure Frame Selector:

   ControlledReceives = Both[7]

   ControlledSends = Untagged

   Secure Frame Verification:

   ValidateReceivedFrames = False

   ReplayProtect = False

   The KaY management parameters are set as follows:

   CipherSuiteSelectable = True only for the Default Cipher Suite

   This step should not break anything, unless there is illicit use of the MACsec EtherType. It would be nice at this step if the KaY provided sufficient controls and recording for its Discovery process to allow connectivity to be confirmed. Some testing of key agreement protocols is also possible at this stage. Both of these suggestions lie outside the scope of this note.

2. For the next step the KaY has to be set so that it will generate a stream of SAKs even if the key agreement protocols are not working properly. Change the SecY management parameters as follows:

   Secure Frame Selector:

   ControlledReceives = Both[8]

   ControlledSends = Tagged

   Secure Frame Verification:

   ValidateReceivedFrames = False

   ReplayProtect = False

   This step shouldn't break anything either. The counts of SecTAG'd frames received should go up, and those of untagged frames should go down. Monitoring those counts together with knowledge of protocols that run over the Uncontrolled Port should be sufficient to confirm that all Controlled Port traffic on the LAN is now being sent tagged.

3. Change the SecY management parameters as follows:

   Secure Frame Selector:

   ControlledReceives = Tagged[9]

   ControlledSends = Tagged

   Secure Frame Verification:

   ValidateReceivedFrames = False

   ReplayProtect = False

   Check that the connectivity is indeed not being disrupted, that the spanning tree has not reconfigured etc. A deadman timer protected change is a useful tool here, as it will recover an inband managed network if connectivity was broken.

   This is a good time to get the KaY and the key agreement protocol infrastructure really

[3] Much more likely when the LAN is "virtual", as in the case of a service instances provided by a Provider Bridged Network.

[4] One of the worst cases is intermittent failure of key agreement protocols due to intermittent failure of the infrastructure components that support them.

[5] As is readily apparent this is just a top-of-mind list. A structured and time tested decomposition of failure causes would be appreciated.

[6] These follow P802.1AE/D2 very loosely. That spec also needs to describe which parameters are expected to be set directly, and which should be set by the Kay.

[7] The SecY verifies frames with the MACsec Ethertype, but transparently receives all other frames. Other possible settings are All (bypasses the SecY, even frames with MACsec EtherType are transparently passed), Untagged, and Tagged.

[8] The SecY verifies frames with the MACsec Ethertype, but transparently receives all other frames. Other possible settings are All (bypasses the SecY, even frames with MACsec EtherType are transparently passed), Untagged, and Tagged.

[9] The SecY verifies frames with the MACsec Ethertype, but transparently receives all other frames. Other possible settings are All (bypasses the SecY, even frames with MACsec EtherType are transparently passed), Untagged, and Tagged.

working. When it is working the InvalidReceivedFrames count should stop incrementing. Check the UnknownSCI and UnknownSC counts if problems persist (Discovery may not be working correctly), if there are none of these check the per SA counters to see if SAs are being used correctly by the transmitter, otherwise suspect the keys. Look at the KaY to monitor key agreement and SAK derivation.

4. Allow the KaY to select other Cipher Suites, including those which set the E bit, i.e. require cipher suite selection, and possibly the SAK, to be known at receivers if connectivity is not to break.

   Confirm that frames are not being discarded as invalid on receipt.

5. Set ValidateReceivedFrames = True. If the previous steps were OK nothing should change. Check the ReplayViolations count. If it is not incrementing then Replay Protect can be set.

6. Check the spanning tree roles and the authorization provided by the KaY against those that would be permitted by RSTP policies. Set the restrictedRole (controlling whether the Port can be a Root Port) and restrictedTCN policies. Verify that the spanning tree configuration has not changed.

Done.

## RSTP policies

Secure RSTP implements policies that use the authorization provided by P802.1af and the associated integrity and origin guarantees provided by MACsec. RSTP policies are applied to received BPDUs and control whether or not:

1. The receiving port can be a Root Port (restrictedRole)
2. Topology changes are accepted (restrictedTcn)

If restrictedRole is set for a port then the RSTP's[10] updateRoles procedure will not select it as a Root Port, but only as a Designated, Alternate, or Backup Port. If a BPDU has been received that would (if restrictedRole were not set) it becomes an Alternate Port instead. This means that the spanning tree configuration of the network 'behind' the bridge port cannot be changed by receipt of a BPDU by the port[11]. It also means that there will be no connectivity through the port if a BPDU conveying a priority vector suggesting a better Root or better path to the Root is received.

If the restrictedTcn parameter is set, no topology changes are propagated through the port to the other ports of a bridge.

## Unauthorized Bridges

It may be the case that there are unauthorized bridges attached to the LAN, and that the intent is to deploy MACsec to excluded these. Since both the MAC Relay and RSTP use the Controlled Port, the unauthorized bridges will be excluded from both the control protocol and the data when ValidateReceivedFrames is turned on.

To minimize the chance of those unauthorized bridges blindly passing MACsec frames and causing a loop in the network, the SecYs for all Ports that are not yet attached to secure LANs and are not currently participating in deployment to a specific LAN should have ControlledReceives = Untagged.

## Staging capabilities

It is useful if deployment plans not only recognize the challenge of upgrading different systems at different times but also the opportunities presented by the staged availability of capabilities within a single system. While these may only provide a fraction of the anticipated benefits of protocol deployment they present an opportunity to debug a large part of the deployment.

However such fractional ability may also require additional controls and different system configurations. Their benefits should be balanced against the complexity of incorporating these in the system.

## Soft MACsec deployment

The term 'soft deployment' is used in this note for the idea that a software based MACsec implementation might be used to protect control protocols, in advanced of the availability of full rate hardware capable of protecting all the data[12]. If anyone can think of a better term for this I would appreciate it.

Soft deployment definitely has its problems, as it threatens to divorce the connectivity provided to the control protocols from that of the data that they are meant to be controlling. For bridging that could be disastrous so special care is required.

What soft deployment can do is to

1. Ensure that the network as a whole is not disrupted by an unauthorized bridge that claims to be the Root Bridge or provide a path to the Root, or that injects unwanted topology change notifications into the rest of the network.
2. Detect bridges that are unauthorized but have not been attached to the network with any malicious intent.

[10] This a proposal not part of the current RSTP specification in 802.1D-2004.

[11] This has to be checked to ensure that no proposal-agreement handshakes can be initiated.

[12] A major point of this note is to describe how this can be done for RSTP. The idea that it could be done was advanced by Norm Finn. I am not sure if the description of how given here fits Norm's ideas on the subject.

What it can’t do is protect data.

The MAC Relay is connected to the Uncontrolled Port.

The RSTP entity is connected to both the Controlled and Uncontrolled Ports. It always transmits and receives frames using the Controlled Port, and also receives BPDUs on the Uncontrolled Port when ControlledReceives is set to Tagged.

If the goal is simply to restrict the impact of BPDUs from unauthorized bridges, BPDUs are also transmitted using the Uncontrolled Port when ControlledSends is set to Tagged.

.A slightly higher priority and subtly different identity are associated with BPDUs received from the Controlled Port, thus ensuring that they, and their associated authorization level, are not immediately displaced by the Uncontrolled BPDU from the same transmitter.

Different levels of authorization are naturally attached to BPDUs received from the Controlled and Uncontrolled Ports. MACsec protected BPDUs are used to establish spanning tree port states for Uncontrolled Port data, the two are not treated as separate Bridge Ports. All the bridges attached to a LAN have to use the same relative priority for BPDUs if data connectivity is provided, so there can be times when the reception of an Uncontrolled BPDU can displace a priority vector received in a Controlled BPDU and cause the receiving port not to be a Root Port. The authorization level associated with each received priority vector is maintained along with the RSTP “infoIs” variable for the port to ensure that the policy control is applied correctly.