FSMLabs Technology White Paper. © FSMLabs 2016.

Best Practices with TimeKeeper® Sourcecheck

The TimeKeeper 'Sourcecheck' feature is an analytics and automated threat detection tool that allows TimeKeeper to detect bad time sources and switch away from them. Bad time can come from misconfiguration, equipment and software failures or intentional time based attacks. Sourcecheck can detect and defeat them.

Sourcecheck Overview

Sourcecheck is a configurable option in TimeKeeper and should only be selected if there are high quality alternate sources.

When source check is disabled, TimeKeeper will accept the highest priority configured source (NTP, PTP, GPS or other) as authoritative. TimeKeeper will apply sophisticated filtering to reduce jitter and it will failover to the next source and produce alerts if the highest priority source stops communicating new times (or violates protocol). If the highest priority source resumes operations later, TimeKeeper without source-check will revert. This behavior avoids a number of serious problems caused by NTPs attempts to second guess time and by the simplistic "Best Master Clock" operation of PTP 1588 while preserving interoperability.

When Sourcecheck is enabled TimeKeeper will cross-check all the existing time sources with one another to validate them. If the highest priority time source fails the Sourcecheck validation TimeKeeper will mark it as invalid, generate alerts and switch to the next highest priority time source that is not currently marked invalid. Cross check detects of bad time from equipment failures, misconfiguration, false leap seconds as well as malicious activities such as GPS spoofing and other time based attacks.

How it works

Sourcecheck combines a multi-source stochastic analysis, protocol verification, and source voting to detect outlier time sources. Each time source gets a vote for the correct time. TimeKeeper then looks for patterns over an interval to see if a second source is more in sync with a critical mass of sources.

To the right you can see a simple example of a tie with 3 time sources. The primary (source 0) and two secondary time sources all disagree. The secondary time sources disagree from the primary by +0.5 and -0.5 milliseconds respectively. Since there is no 'quorum' here source 0 remains valid.

Sourcecheck tie

Here secondary time sources 1 and 2 do not agree so all sources remain valid

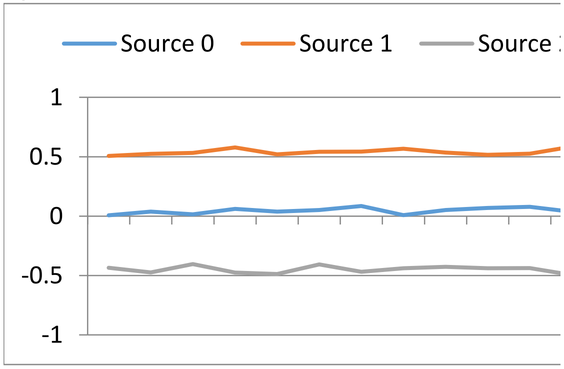

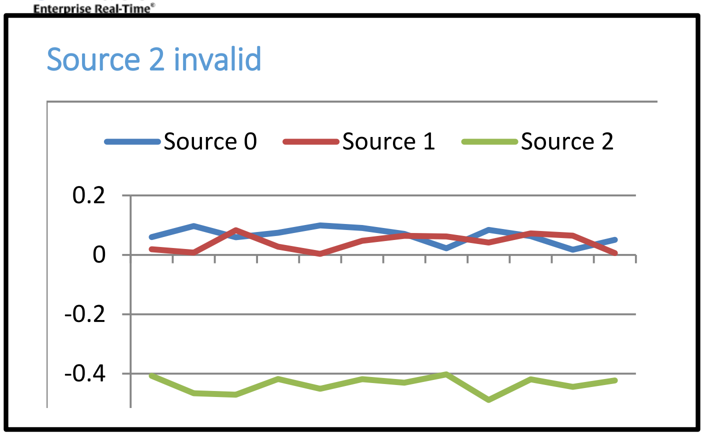

To the left you can see a chart showing that source 0 and 1 agree while source 2 is offset by -0.5 milliseconds. Sourcecheck will declare source 2 as invalid in this case.

Below you can see an example where source 0 would be declared invalid. That is because source 1 and 2 form a quorum.

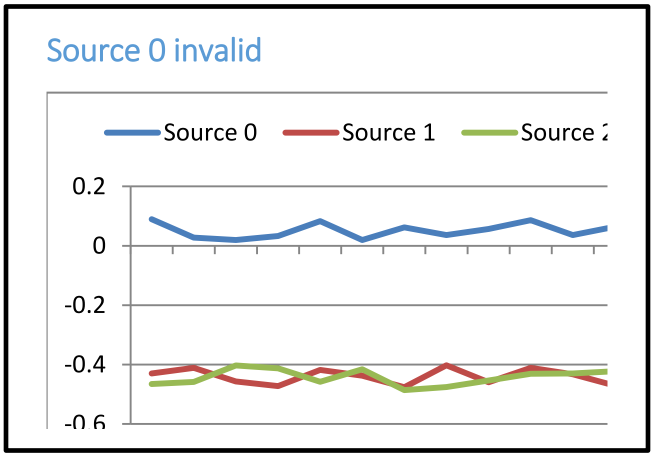

Sourcecheck analyzes both time and frequency. Frequency analysis can be especially effective in a GPS spoof attack. Note the graph below.. It shows the two secondary time sources halfway through start slowly drifting in offset because a GPS spoof attack is under way which is slowly moving the time on our primary time source (source 0). In this graph it looks as though all the other time sources are moving. Eventually our time offset analysis would catch the compromise but frequency analysis spots the problem earlier.

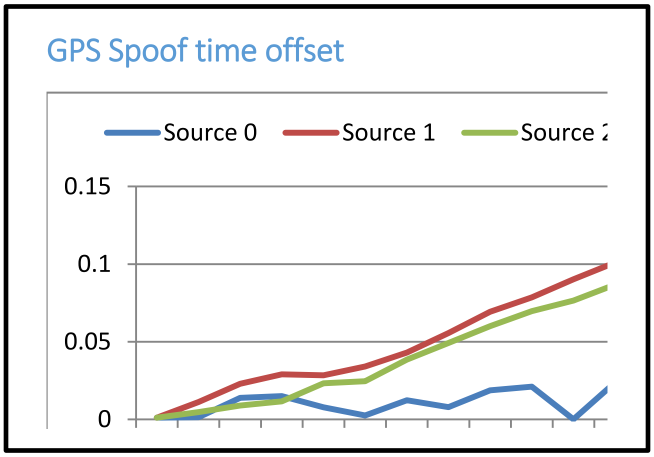

The graph below shows the rate of change (frequency) for each time source. It's quite clear from this graph that TimeKeeper will quickly catch this difference in frequency with Sourcecheck and invalidate the primary source which is being spoofed even though it's not initially obvious in the time offset graph.

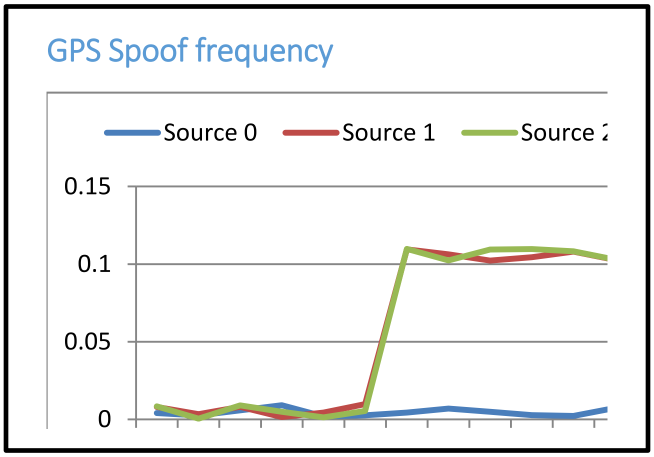

Restoring a previously invalid source

Once a time source corrects its time and begins passing Sourcecheck validation it will not be marked 'valid' immediately. Sourcecheck will wait 15 minutes before allowing a source to be considered valid. This avoids switching back to a bad time source because it momentarily shows good time or quickly switching between time sources that are in the 'gray area' of being marked invalid.

Leapsecond checks

In addition to the cluster analysis that Sourcecheck performs it will execute additional checks. To avoid false leap seconds which have plagued other solutions Sourcecheck will confirm that a majority of time sources show a leap second before accepting the leap second as valid. Sourcecheck will also only permit leap seconds during certain times of the year during which leap seconds are inserted. If a leap second is advertised by a time source outside of these ranges then it is rejected.

Some configurations can result in unexpected behavior so it's important to be aware of that. If a primary time source is configured to 'slew' in response to a leapsecond but backup time sources are configured to 'step'. In this case immediately after the leap second Sourcecheck would see the primary time source smoothly/slowly slewing the time to correct for the leap second while the backup time sources stepped immediately. That would cause Sourcecheck to reject the primary time source since it disagrees with the two backup time sources. After the primary time source completes the slew it would be accepted by Sourcecheck again as all time sources then agree.

How to setup Sourcecheck properly

The most critical aspect of setting up Sourcecheck is to select good quality time sources. If you select multiple time sources that are outside of your control, are unreliable and cannot be trusted then you have given these sources the ability to invalidate your primary source. So it is always recommended that you use only your own corporate or internal time sources (no public NTP servers). More is always better in this case since Sourcecheck can do some cross-checks with only two sources but three or more are necessary for full validation. In general we recommend 5 sources for a proper configuration.

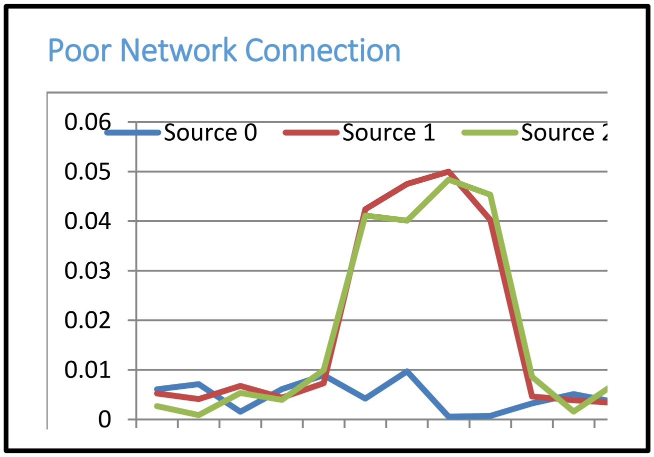

When all time sources are within a certain threshold (configurable) then they are all declared to be 'valid' and Sourcecheck will not invalidate any of them. The Sourcecheck threshold setting avoids spurious source changes where all time sources are accurate within a reasonable margin of error (perhaps a few hundred nanoseconds on a decent network). In general it's not necessary to adjust the default value (30 microseconds) but it can be increased to allow inclusion of lower quality sources in the cross-check. For example, a known-high jitter network link or a poor quality time source that has a known quality can be used in with Sourcecheck by adjusting the threshold.

A good practice when configuring Sourcecheck for the first time or with new sources is to first setup the candidate time sources without Sourcecheck enabled. Over a day or two, TimeKeeper will collect data on how the time sources behave. The graph to the right shows source 0 as a local time source. Sources 1 and 2 are remote time servers. During this test the remote time sources show a large offset for a while which is likely due to a network disruption. If Sourcecheck had been enabled it would have declared source 0 as invalid since it disagrees with sources 1 and 2. These cross-check sources would not be good candidates here without additional cross-check sources that do not use this likely unstable network connection.

Avoiding loops

Something that must be avoided when configuring Sourcecheck is cycles or loops in a time network. For example, one can configure 3 different systems to cross-check with all other systems. It's possible for a situation to occur where the primary time source is rejected and a backup time source is used on all of these systems. In that case, A might track B, B might track C and C might track A. No system is tracking a valid time but they are all tracking one another which can result in time being incorrect.

How to purchase

TimeKeeper, TimeKeeper Server Software, and TimeKeeper Client Software are all available from FSMLabs and FSMLabs' resellers. For purchase information or for a live demonstration of TimeKeeper please contact FSMLabs at sales@fsmlabs.com.