

Paper No. 990448

**An Effectiveness Analysis of SafeStat
(Motor Carrier Safety Status Measurement System)**

November 1998

By

David G. Madsen

US DOT - Volpe National Transportation Systems Center

DTS-42

55 Broadway

Cambridge MA 02142

617-494-2030

Fax 617-494-2787

madsen@volpe.dot.gov

Donald G. Wright

US DOT - Volpe National Transportation Systems Center

DTS-42

55 Broadway

Cambridge MA 02142

617-494-2751

Fax 617-494-2787

wrightd@volpe.dot.gov

An Effectiveness Analysis of SafeStat (Motor Carrier Safety Status Measurement System)

ABSTRACT

SafeStat (Safety Status Measurement System) is an automated analysis system developed for the Federal Highway Administration's (FHWA's) Office of Motor Carriers (OMC) that combines current and historical safety data to measure the relative safety fitness of interstate motor carriers. SafeStat enables the OMC to quantify and monitor the safety status of motor carriers and guides the deployment of resources to focus on carriers posing the greatest safety risk.

SafeStat resulted from research performed by the U.S. DOT's Volpe Center for the OMC to improve motor carrier safety fitness assessment and prescribe actions to correct safety deficiencies. SafeStat was initially developed and implemented as part of a federal/state pilot program. It has since been implemented nationally by the OMC to identify and prioritize individual motor carriers for on-site safety compliance reviews.

An effectiveness study was devised to confirm that SafeStat-identified carriers were indeed high safety-risk carriers. This study examines post-identification carrier crash experience and tests SafeStat's effectiveness by comparing the crash rates of identified and non-identified carriers.

The effectiveness study shows that SafeStat identifies carriers likely to have significantly higher crash rates than carriers not identified. Specifically, the carriers designated in the highest "at-risk" category by SafeStat had a much higher crash rate (169% greater) following their identification than carriers not identified as safety risks. In addition to confirming SafeStat's effectiveness, the results of the study enable SafeStat developers and the OMC to assess the relative strengths of SafeStat's component parts and to continue to make enhancements to further improve its efficiency.

SAFESTAT DEVELOPMENT

SafeStat (Safety Status Measurement System) is an automated analysis system developed for the Federal Highway Administration's (FHWA's) Office of Motor Carriers (OMC) that combines current on-road safety performance information, the results of on-site safety audits (known as compliance reviews), and enforcement history to measure the relative safety fitness of interstate motor carriers. The system provides the OMC with the capability to quantify and monitor the safety status of motor carriers, especially unsafe carriers, on a continuous basis. The objective of SafeStat is to guide the OMC enforcement and education programs by providing a basis for the efficient allocation of resources to focus on carriers that pose the highest risk of crash involvement. SafeStat was developed as part of the OMC's ongoing effort to be a performance-oriented and analysis-driven organization, and takes full advantage of improvements in safety related data reporting and information technology.

SafeStat was conceived through research performed at the U.S. Department of Transportation's John A. Volpe National Transportation Systems Center (the Volpe Center) in Cambridge, MA. The purpose of this research was to define and propose an improved process to assess motor carrier safety fitness and prescribe appropriate OMC actions when safety deficiencies are detected. SafeStat was initially developed and implemented as part of the Commercial Vehicle Information System (CVIS), now the Performance & Registration Information Systems Management (PRISM) federal/state pilot program. As part of CVIS/PRISM, the Volpe Center developed SafeStat (1) to identify and monitor motor carriers for an experimental safety improvement process, the Motor Carrier Safety Improvement Process (MCSIP).

After proving its usefulness in the CVIS/PRISM pilot program, SafeStat was implemented nationwide by the OMC in March of 1997 to identify and prioritize carriers for on-site compliance reviews. SafeStat continues in use for this purpose with results being updated every six months. Improvements to SafeStat are implemented in new versions with each six-month cycle.

SAFESTAT METHODOLOGY SUMMARY

SafeStat evaluates the relative safety status of individual motor carriers with respect to the rest of the motor carrier population in four analytic Safety Evaluation Areas (SEAs) - Accident, Driver, Vehicle,

and Safety Management. SafeStat uses up to 30 months of motor carrier safety and normalizing data to develop measures and indicators in the four SEAs. The four SEA values are then combined into an overall safety status assessment known as a SafeStat Score.

SafeStat determines a value for each SEA in which a carrier has sufficient safety data. Each SEA value approximates the carrier’s percentile rank relative to all of the other carriers that have sufficient data to be assessed within the SEA. SEA values range from 0 to 100. The higher a carrier’s SEA value, the worse its safety status. An unacceptable SEA value is defined as any SEA value in the worst 25th percentile (i.e., a value of 75 to 100).

SafeStat only identifies carriers with sufficient safety event data that have the poorest safety status. Specifically, SafeStat produces a SafeStat Score for each carrier found with two or more unacceptable SEA values. SafeStat further delineates the worst of these scored carriers as "at-risk." An "at-risk" carrier must be unacceptable in three or more SEAs, with an unacceptable Accident SEA counting twice. This approach is designed to identify the carriers with the worst safety performance at any given time, and hence, the most logical candidates for safety improvement programs or enforcement action.

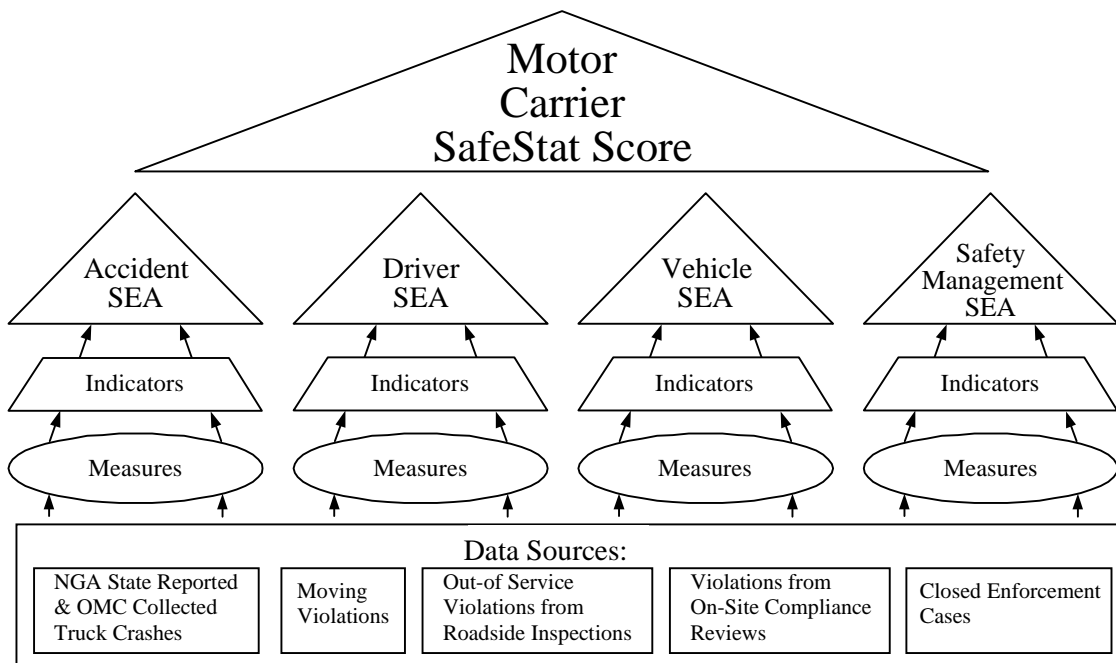


FIGURE 1. Summary of SafeStat Methodology

SAFESTAT EFFECTIVENESS ANALYSIS

As part of the evaluation of CVIS/PRISM, an effectiveness study was devised to confirm that the carriers that SafeStat was identifying were indeed high safety risk carriers. Safety risk at any given time is defined as the likelihood of having crashes in the near future. By examining the SafeStat post-identification crash experience of identified carriers, this study essentially tests SafeStat's crash rate prediction capability and represents the "bottom-line" assessment of its performance. Beyond confirming SafeStat's effectiveness, the results of this study are being used to refine SafeStat to further emphasize the components of the system that are the most closely related to high future crash rates and to evaluate the contribution of potential new measures and indicators.

DESCRIPTION OF THE STUDY

The effectiveness study was accomplished by: (1) performing a simulated SafeStat carrier identification using historical data; (2) observing the crash involvement over the immediate 18 months after SafeStat was run for both the carriers identified by SafeStat as having a poor safety status and other carriers not so identified by SafeStat, but which had sufficient data to be identified; and (3) comparing the post-identification crash rates of both groups of carriers. If SafeStat is effective in identifying unsafe carriers (i.e., carriers having a high risk of being involved in future crashes), then the carriers identified as having a poor safety status would be expected to have higher post-selection crash rates than the carriers that were not identified by SafeStat. The greater the post-selection crash rate for the identified carriers relative to those carriers not identified, the more effective SafeStat would be in identifying unsafe motor carriers.

Rather than use the most recent available data and having to wait for a period of time to collect post-identification crash data, the analysis was performed using historical data. The study was conducted by simulating a carrier identification by SafeStat on data available at an earlier date (April 1, 1996) and then observing the carriers' crash involvement that occurred over the next 18 months (from April 1996 to October 1997). This procedure simulated carrier identification by SafeStat as if it had been run as of April 1, 1996 using safety events that occurred prior to that date, and allowed for sufficient subsequent crash reporting to accurately measure the post-identification crash rates.

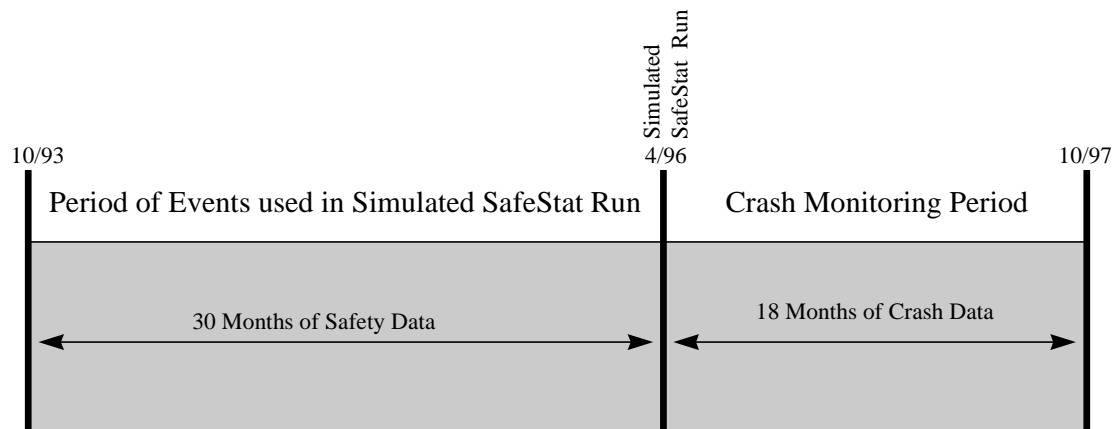


FIGURE 2: Effectiveness Analysis Timeline

From this simulation run of SafeStat, carriers that had sufficient data to be scored were placed into the following groups based on their overall SafeStat results in order to compare the “post-selection crash performance”:

1. carriers identified as “at-risk” (worst SafeStat Scores – carriers that have three or more unacceptable SEAs with the unacceptable Accident SEA counting twice),
2. other carriers identified as having a poor safety status according to SafeStat (carriers that have two unacceptable SEA without an unacceptable Accident SEA), and
3. carriers with sufficient data but not identified by SafeStat as having a poor safety status.

The post-identification crash rate of each group was calculated as the number of reported crashes per 1000 power units (PUs). The number of PUs is defined by the total number of trucks, tractors, hazardous material tank trucks, motor coaches, school buses, minibuses/vans, and limousines that are owned or term leased by a motor carrier. The carrier PU information was based on census data that reside in the centralized OMC national database, the Motor Carrier Management Information System (MCMIS).

The crash data were based upon the crashes reported by the states (according to the National Governors’ Association (NGA) standard) that occurred during the post-selection period (April 1996 to October 1997). These data also reside in the MCMIS. Each reported crash was weighted based on the severity and timing of the crash.

The severity-weighting scheme placed emphasis on crashes with greater consequences, while the time weighting placed emphasis on crashes that occurred soon after the SafeStat identification run. Severity weights were assigned as follows: a weight of 0.5 for property damage only, a weight of 1.0 for crashes involving injuries/fatalities or hazardous material release, and a weight of 1.5 for crashes involving injuries/fatalities and hazardous material release. Time weights were assigned to crashes as follows: a weight of 1.5 for crashes that occurred during the first six months after the SafeStat run, a weight of 1.0 for crashes that occurred 7 to 12 after the SafeStat run, and a weight of 0.5 for crashes that occurred 13 to 18 months following the SafeStat run. Each crash had its severity weight multiplied by its time weight to obtain an overall weight. In each carrier group, the weighted crashes were summed and divided by the number of PUs to provide a weighted crash rate for the group.

DATA ISSUES

The number of PUs was chosen as means of measuring exposure to normalize the state-reported NGA crash data. Assuming relatively consistent vehicle utilization rates, the number of PUs provides an estimate of the amount of time spent traveling when crashes can potentially occur. Discussions with other developers of safety measurement systems (Ontario and Quebec) supported the use of power units to calculate crash rates. Vehicle miles traveled (VMT) data, another popular measurement of exposure, were not used because such data were either not available or current for most carriers. (VMT data as a measurement of exposure also has potential bias problems of overstating the exposure and hence favoring long-haul carriers (primarily operating at high speeds on interstate highways) relative to carriers with operations of short-hauls (primarily at low speeds, often on local roads)).

There is a concern that inaccurate PU data (especially in cases where the PUs are understated) can bias the results of the effectiveness study. To mitigate this potential bias due to inaccurate PU information, the Poisson distribution was used to identify carriers that had unreasonably high crash rates for the post-identification period. While vast majority of the carriers (74,073) had reasonable crash rates, 52 carriers were identified as having unreasonably high crash rates, which were probably based on inaccurate PU normalizing data. Thus, these carriers were not included in the effectiveness study results.

State-reported NGA crash data were used in the study because these data comprise the largest/most-complete set of carrier crash information that can be linked to the specific carriers involved. While the NGA data does not provide a complete record of all carrier crashes, a recent OMC analysis estimated that majority of all interstate carrier crashes are being recorded in the NGA data. There may be concern that the missing crash information could possibly bias the results of the effectiveness study. But the likelihood of a crash being recorded or not in the NGA data is independent of whether the carrier has been identified by SafeStat. This independence makes the NGA data a large, unbiased sampling of crashes to be used in the post-selection period. Another potential problem of using NGA crash data is the potential for long delays between when the crashes occur and when the crashes are recorded in the NGA data. A vast majority of the crashes, however, are entered into the data within 6 months. The effectiveness study used crash data available as of April 1998, six months after the end of the post-selection monitoring period. The use of this cutoff date ensured that most of crashes that occurred during the monitoring period and were eventually recorded in the data were used in the study.

RESULTS

Overall Effectiveness of SafeStat

The post-selection crash rates for the SafeStat identified and not identified carrier groups were examined both in terms of their overall SafeStat Scores and in terms of the four Safety Evaluation Areas (SEAs) — Accident, Driver, Vehicle, and Safety Management — that determine the overall SafeStat Scores.

TABLE 1. Post-Selection Crash Rates

Carrier Group	Number of Carriers	Weighted Crash Rate*	% Higher than Not Identified Carriers
All Identified	4,276	56.4	85%
At-Risk (with Worst SafeStat Scores)	1,450	82.3	169%
Other Identified (with Poor SafeStat Scores)	2,826	43.2	41%
Not Identified	69,797	30.5	-

* Weighted crashes per 1000 power units

These results confirm that SafeStat did identify carriers with a higher crash risk. The group of all carriers that SafeStat identified as poor performers had an 85% higher crash rate than carriers that were not identified. The carriers designated as “at-risk” by SafeStat had a much higher crash rate (169% greater) than the carriers that were not identified. A majority of these “at-risk” carriers were identified in part because they had previous problems with respect to their crash rates (i.e., they had unacceptable Accident SEA values).

However, even the SafeStat identified carriers in the “other identified” group, which did not have high Accident SEA values but were in the worst 25th percentile in two of the other SEAs, posed a 41% greater crash risk than the carriers that were not identified. This result shows that SafeStat has the proactive capability to identify carriers that are likely to be involved in crashes even though they previously did not have exceptionally high crash rates.




FIGURE 3. Crash Rates for the Three Groups of Carriers

Effectiveness of Individual SEAs

Further testing was done to determine the effectiveness of the principal components of SafeStat. This was accomplished by placing carriers into groups based on their performance results for each particular SEA (i.e., Accident, Driver, Vehicle, or Safety Management).

The results for carriers with high individual SEA values compared to those with lower SEA values are as follows (Carriers with high SEA values were in the worst 25th percentile and were designated as the

worst performers in that particular evaluation area. Conversely, carriers with no high SEA values were not in the worst 25th percentile, and therefore, were not among the poorest performers in that SEA.):

TABLE 2. Crash Rates of Carriers with and without High SEAs

Safety Evaluation Area	Number of Carriers	Weighted Crash Rate*	% Greater than Carriers without the High SEA
High Accident SEA	2,596	81.4	172%
No High Accident SEA	71,477	29.9	-
High Driver SEA	7,036	56.2	90%
No High Driver SEA	67,037	29.5	-
High Vehicle SEA	12,456	38.3	22%
No High Vehicle SEA	61,617	31.4	-
High Safety Mgmt. SEA	4,442	42.0	35%
No High Safety Mgmt. SEA	69,631	31.0	-

* Weighted crashes per 1000 power units from 3/96 to 10/97

Accident SEA

The results confirm what may seem intuitively to be obvious: carriers with high crash rates in the past are likely to continue to have high crash rates in the future. In other words, past crash rate performance is a good indicator of future crash rate performance. The effectiveness study shows a 172% greater post-selection crash rate for carriers with poor Accident SEAs compared to carriers that were not identified as having poor Accident SEAs. Comparing SEAs, the Accident SEA is by far the most effective SEA for identifying high-risk carriers, thereby justifying the “double-weighting” of the Accident SEA in calculating a SafeStat Score.

Driver SEA

The Driver SEA (with a 90% higher crash rate for carriers with poor Driver SEAs) is the next most effective SEA. These results from the study are especially impressive because the criteria for the Driver SEA are based on violations and are independent of crash history.

Vehicle SEA

Carriers with poor Vehicle SEAs did have a higher crash rate (22%) than carriers without poor Vehicle SEAs. Although the difference is not as great as the crash rate differences in the Accident and Driver SEAs, it is significant. As with the Driver SEA, the criteria for the Vehicle SEA are based on violations and are independent of crash history. Also, it should be noted that due to the large amount of vehicle roadside inspection data, Vehicle SEA values were computed for many more carriers than were Accident or Driver SEA values (12,456 as opposed to the Accident SEA's 2,596 and the Driver SEA's 7,036) and, thus, it has the potential of identifying more carriers in absolute terms.

Safety Management SEA

The Safety Management SEA is also effective in identifying carriers with high crash rates. Indicators in this SEA are based on safety regulation compliance supporting the association of safety regulations with crash risk. Carriers with high Safety Management SEAs had a 35% higher post-identification crash rate than carriers that did not have high Safety Management SEAs. Recent improvements made to this SEA in SafeStat have substantially increased its effectiveness.

CONCLUSION

SafeStat does work. The effectiveness study shows that all of the individual parts of SafeStat and SafeStat as a whole do indeed identify carriers that are likely to have significantly higher crash rates than carriers not identified. The effectiveness study has also proven to be a useful tool in quantifying the performance of SafeStat. SafeStat was designed to be continuously improved. The results of the study enable SafeStat developers and the OMC to assess the relative strengths of SafeStat's component parts and to continue to make enhancements to improve its efficiency. Finally, SafeStat continues to be strengthened and improved through the addition of better data and new indicators (most recently, a Moving Violation Indicator in the Driver SEA, which a separate analysis has shown will further increase SafeStat's effectiveness).

REFERENCES

1. A complete description of SafeStat is found in "SafeStat Motor Carrier Safety Status Measurement System Methodology: Version 6," July 1998, The Volpe Center, DTS-42, 55 Broadway, Cambridge, MA 02142.