

Main Ideas and Critiques

The authors tested the ability of operators with different skill sets and levels to perform predefined tasks and solve “injected” problems in a system. All the actions by each operator were recorded in a trace for later playback; the mistakes each operator made were categorized by their effects and the root type of the mistake. The most common mistakes were misconfigurations of applications, and degraded throughput was the most common effect of mistakes. An increased skill level of the operator did not correlate to fewer mistakes, though the experiment design did not balance the difficult tasks across all skill levels.

The paper presents the results of a prototype system where two network slices are maintained on the same network, connected by one-way shunts of data flow. The online slice is the production system, and a validation slice is used as a temporary location for a system with a new configuration to be tested.

Two types of validation were used and analyzed—replica-based and trace-based validation. Replica-based validation involves the synchronization of state between an online server and the validation server, and subsequent real-time comparisons of inputs and outputs between the online and validation servers. One issue here is that, since the validation slice cannot affect the online slice, when a significant deviation between the two servers’ states appears, there is no way to continue validation since no reputable comparison can determine whether deviations in input/output are due to the difference in state or a problem with the configuration of the server being validated.

The second type of validation is trace-based validation, where the entire input/output requests related to a server to be validated are stored as they are received. At a later date, a validation server can be passed the relevant requests, and its results can be verified.

The paper raised an issue regarding the age of traces—in the described trace-based validation system, the server being validated has access to the lower tiers of the system; when the traces are old, the state of the underlying system may change, resulting in different results. I don’t understand how the server being validated could have access to lower levels of the system—it seems that the responses from the lower levels of the system should have been stored as well during the trace-recording stage—if this were the case, the current state of the online system (i.e. the database) would not be an issue regardless of the age of the traces as the trace data would contain the snapshot of the underlying online system.

One final issue that I have with the Mendosus slice isolation tool is that it fails to take into account the effect of excess network traffic causing congestion in the network. Though this is an unlikely scenario, it is possible that a server undergoing validation might send large amounts of packets into the network that, although not targeted towards servers in the online slice (and hence not blocked by Mendosus), would still affect the network traffic in the online slice. A stronger architecture for the Mendosus online-validation segmentation would involve physical network segregation with the exception of the shunts between the two slices of the network.