

Calculating the Size of the Nonpersistent Event Queue

If communication between the SAE and VTA is lost for an extended period, a large backlog of events can build up in the SAE fail queue. This backlog can result in the SAE sending events far in excess of the average rate for an extended period. Note, however, that smaller event queue sizes affect recovery time if the connection between the SAE and the VTA is disrupted. When the VTA event queue fills, the VTA signals the SAE to stop sending events, and the SAE does not send any events for 60 seconds. If the VTA event queue is limited to less than the number of events the VTA can process in 60 seconds, it will empty the queue and be idle until the SAE starts sending it events again.

For example, if you set the VTA event queue size to the number of events the VTA can process in 30 seconds, after an SAE-VTA communication disruption the SAE rapidly sends events that it has stored in its local fail queue. After the VTA has collected the events that take it 30 seconds to process, it signals the SAE to stop sending events, and the SAE stops for 60 seconds. The VTA will complete processing the events in its queue in 30 seconds and then will be idle for 30 further seconds before the SAE starts to send events again. In this example, the VTA could clear the backlog of events in the SAE's fail queue twice as fast if the VTA's event queue was large enough to hold the number of events it can process in 60 seconds.

Contact the Juniper Technical Assistance Center or Juniper Professional Services if you need further advice about sizing the VTA event queue.

