

Setting the Size of the Event Queue

The VTA has an event queue that holds plug-in events from the SAE until the VTA processes them. There are two types of event queues:

- With a persistent event queue, no events are lost even if the VTA or the J2EE application server fails and is restarted. If you are using a persistent event queue, set the size of the VTA event queue to about the number of events your VTA can process in 60 seconds.
- With a nonpersistent (in memory) event queue, events can be lost if the VTA or the J2EE application server fails. However, for performance reasons you may want to configure a nonpersistent event queue.

If you configure a nonpersistent event queue, the event queue size is the maximum number of events that can be lost if the application server or the VTA fails. If the SAE sends events faster than the VTA can process them, the event queue fills, and the VTA signals the SAE to stop sending events for 60 seconds. Your VTA instances should be deployed on a host or cluster with sufficient throughput to handle events above the average rate generated by the SAEs in normal operation. This way, the event queue is sufficient to buffer peak event generation rates.

