

## 4.3 Serverless Computing

An interesting recent trend in cloud computing is the introduction of **serverless computing** as a new paradigm for service delivery. As we show in the chapters ahead, computation and data analysis can be deployed in the cloud via a range of special services. In the majority of cases, the user must deploy VMs, either directly or indirectly, to support these capabilities. Doing so takes time, and the user is responsible for deleting the VMs when they are no longer needed. At times, however, this overhead is not acceptable, such as when you want an action to take place in response to a relatively rare event. The cost of keeping a VM running

continuously so that a program can wait for the event may be unacceptably high.

Serverless computing is similar to the old Unix concepts of a daemon and cron jobs, whereby a program is managed by the operating system and is executed only when specific conditions arise. In serverless computing, the user provides a simple function to be executed, again under certain conditions. For example, the user may wish to perform some bookkeeping when a new file is created in a cloud repository or to receive a notification when an important event occurs. The cloud provider keeps a set of machines running to execute these functions on the user's behalf; the user is charged only for the execution of the task, not for maintaining the servers. We return to this topic in chapters 9 and 18.