Deploying CloudSuite Simics images in your system:

In the CloudSuite Simics release, we distribute a single Simics checkpoint, named phase_000, for each workload, which was saved at the steady state of the corresponding workload's execution.

Before you start running the CloudSuite Simics images, you need to manually configure the absolute checkpoint paths in every workload's phase_000 checkpoint file to reflect your setup's configuration.

The CloudSuite release package contains:
1. An OS and a binary disk, which are commonly used by all workloads (under binary_disk).
2. One or two data disks for each workload (under datasets).
3. Workload image disk.

In the CloudSuite release, each of these logical components is a file. The absolute paths to these files are specified in the checkpoint path list in the phase_000 checkpoint file.
For example, in the Media Streaming workload's phase_000 checkpoint file, you will see three paths specified:

```
checkpoint_path: (         "/path/to/os_disk",
                           "/path/to/binary_disk",
                           "/path/to/media_streaming_data",
"/path/to/cloudsuite/images/streaming/1cpu")
```

All you need to do is to replace these paths with the absolute paths in your setup and then the Simics images will be ready to run. Make sure that you leave the first path empty.

Directory Structure for Simics Checkpoints and Flex-points:
You need to maintain a specific directory hierarchy to store intermediate Simics checkpoints and flex-points.
The recommended directory hierarchy is as follows (under the path *ckptdir* specified in *global.run_job.rc.tcl*):
- Every workload has a root directory.
- Under the workload root directory, there is a directory for different CPU counts of the corresponding workload.
- Under each *CPU* directory, there is a *baseline* directory (corresponds to baseline in .run_job.rc.tcl). You can add other directories under a *CPU* directory, if you decide to generate a new set of phases, then you will need to modify .run_job.rc.tcl accordingly.
- Under *baseline*, there are N phase directories, each containing the corresponding Simics checkpoint under the *simics* subdirectory and M *flexpoint* directories. Flexus expects to find a Simics checkpoint that has the same name as the phase directory, under every phase's Simics subdirectory. For example, the simics directory under phase_000, should contain a Simics checkpoint named phase_000. In our release, we provide phase_000 and the rest of the initial phase checkpoints should be created by the users, depending on the desired degree of simulation parallelism.

- In the release, you can find phase_000 under workloads/workload_name/N_cpu/baseline/phase_000/simics/, named specifying the configuration of the target machine. Under every phase_000/simics directory, the file without any extension is the initial checkpoint of the workload (e.g., under workloads/streaming/1cpu/baseline/phase_000/simics, 1s-4gb-1c-4gb is the initial checkpoint). One should load this checkpoint using Simics (), and save a checkpoint under the same directory using Simics, naming the checkpoint phase_000. After creating the phase_000 checkpoint, the workload is ready to run.

  To load the checkpoint in the example above named 1s-4gb-1c-4gb and save the checkpoint with the name phase_000, please run the following three commands:
  ```
  start-simics
  simics> read-configuration
  workloads/streaming/1cpu/baseline/phase_000/simics/1s-4gb-1c-4gb
  simics> write-configuration
  workloads/streaming/1cpu/baseline/phase_000/simics/phase_000
  ```

- Each *flexpoint* directory contains a *simics* subdirectory to store its corresponding Simics checkpoint. Every user has their own subdirectories where they can store Flexus-generated flex-points under specific subdirectories, which will be named by the user while generating flex-points with Flexus.

In the CloudSuite release, you will find a checkpoint directory called workloads. It contains the directory structure described above,