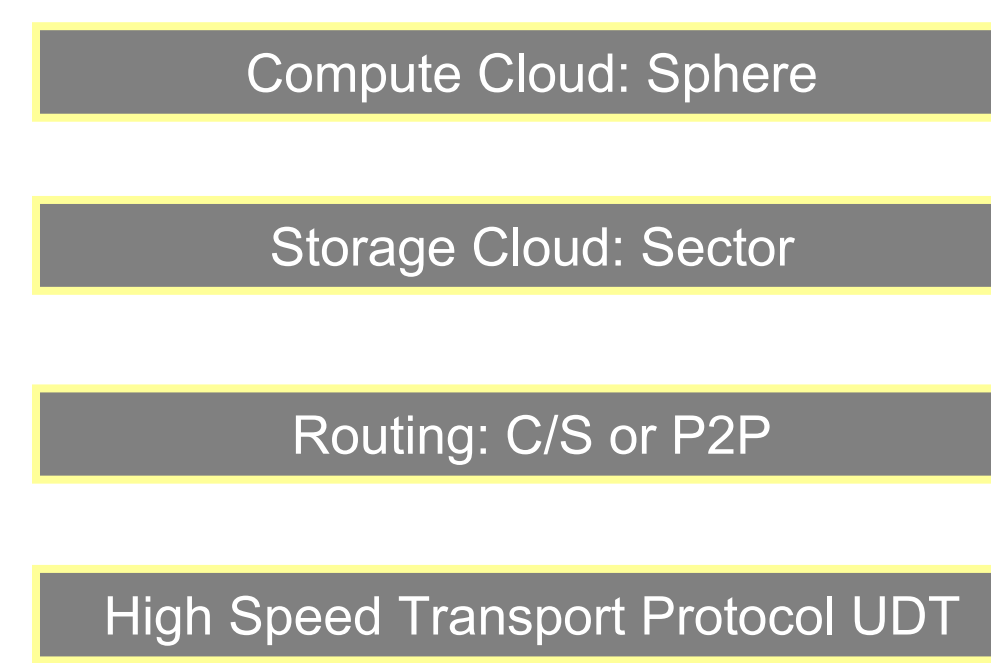
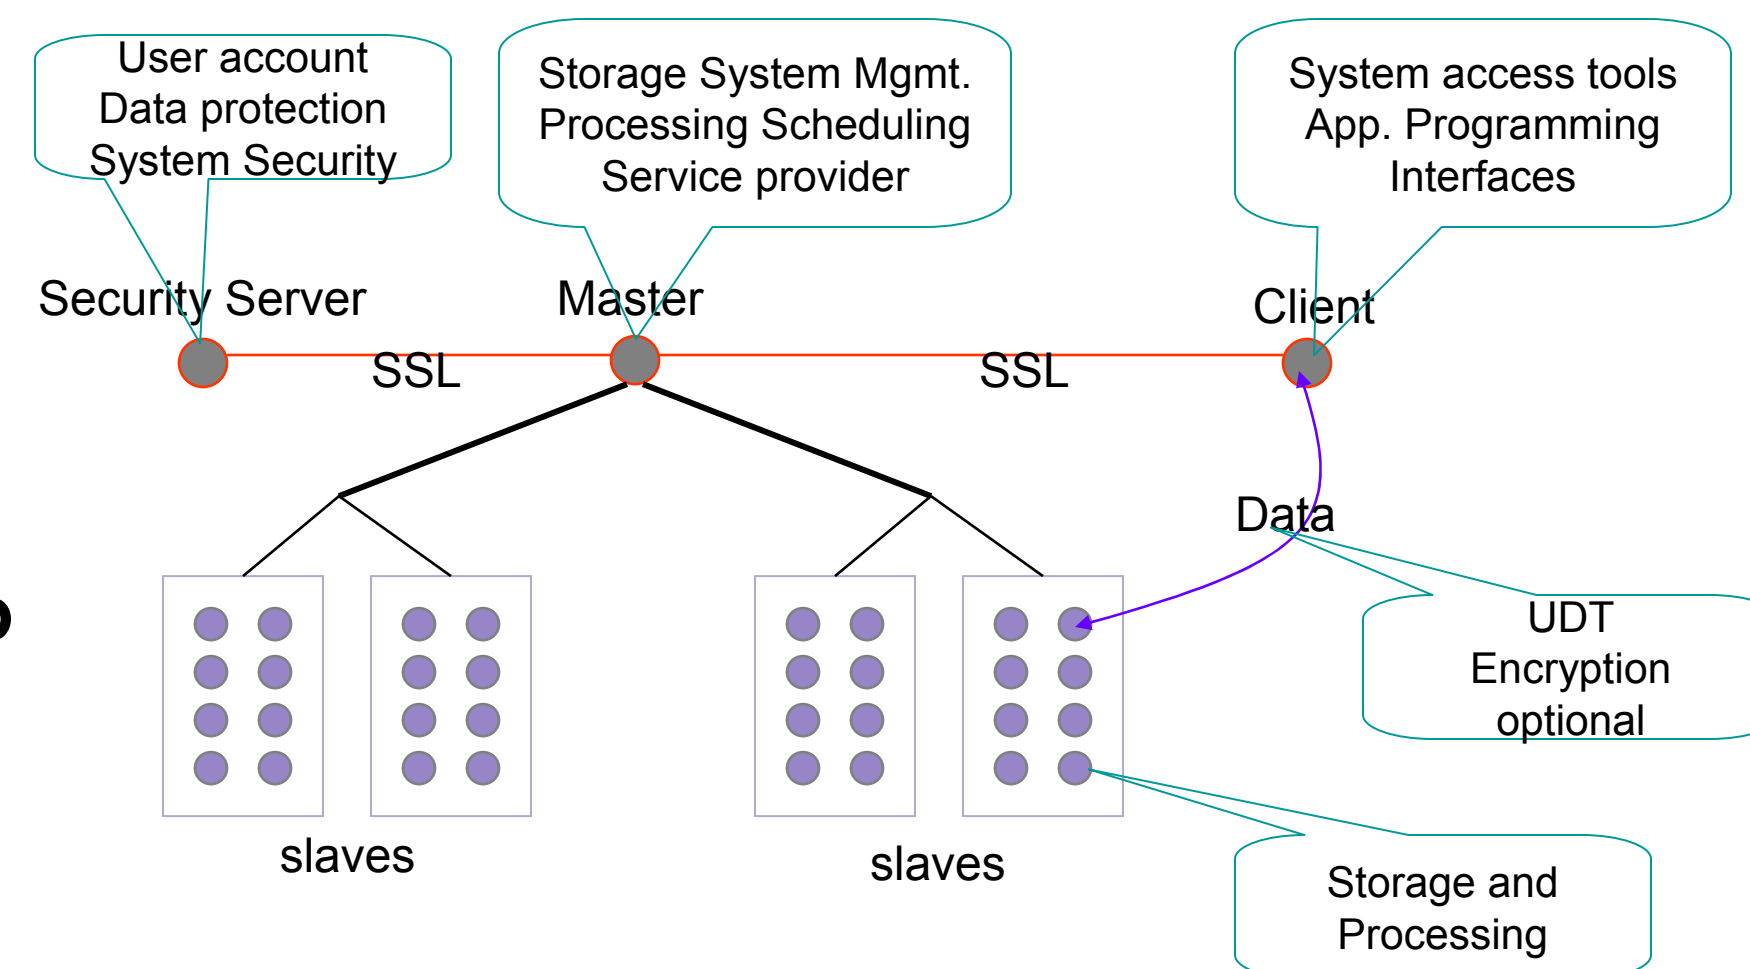


Sector/Sphere is a software platform that supports very large distributed data storage and simplified distributed data processing. The system consists of Sector, a distributed storage system, and Sphere, a runtime middleware to support simplified development of distributed data processing.



In the **Sector** system, a master node maintains the file system, while the data is stored on the slave nodes, possibly across multiple data centers. A security server provides user account verification, access control IP list, etc. UDT is used for high speed data transfer between slaves and between slaves and clients.



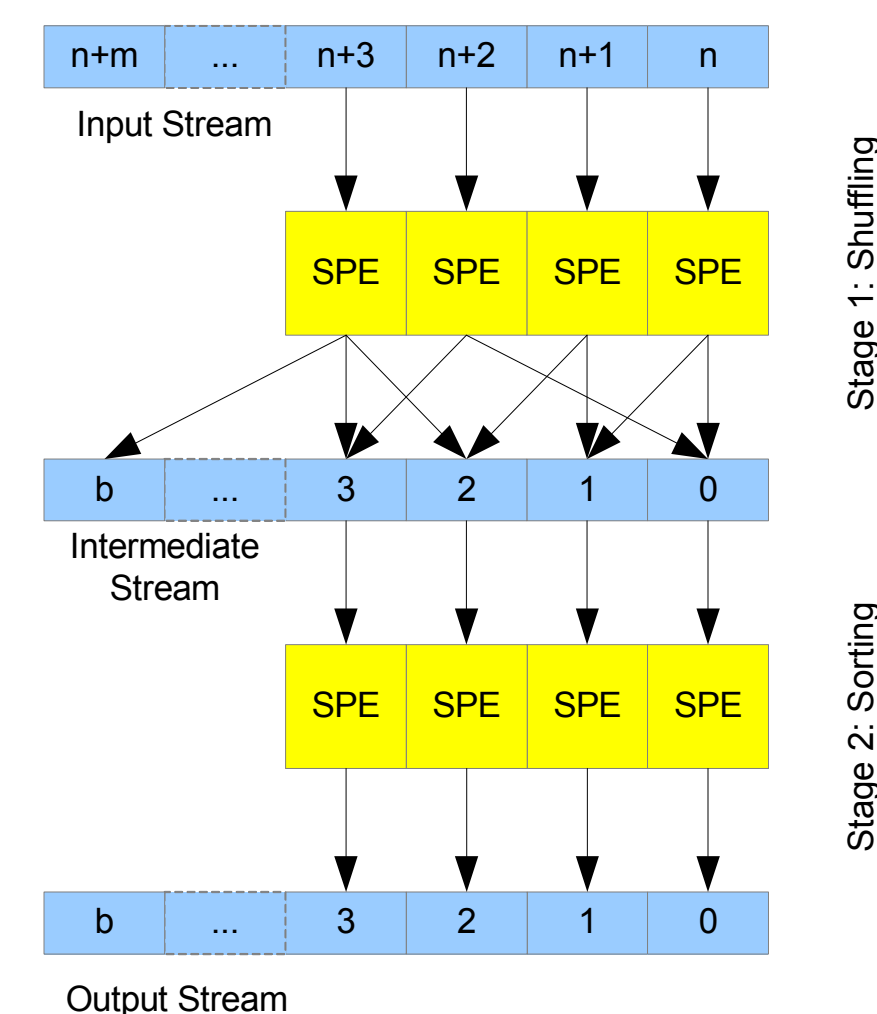
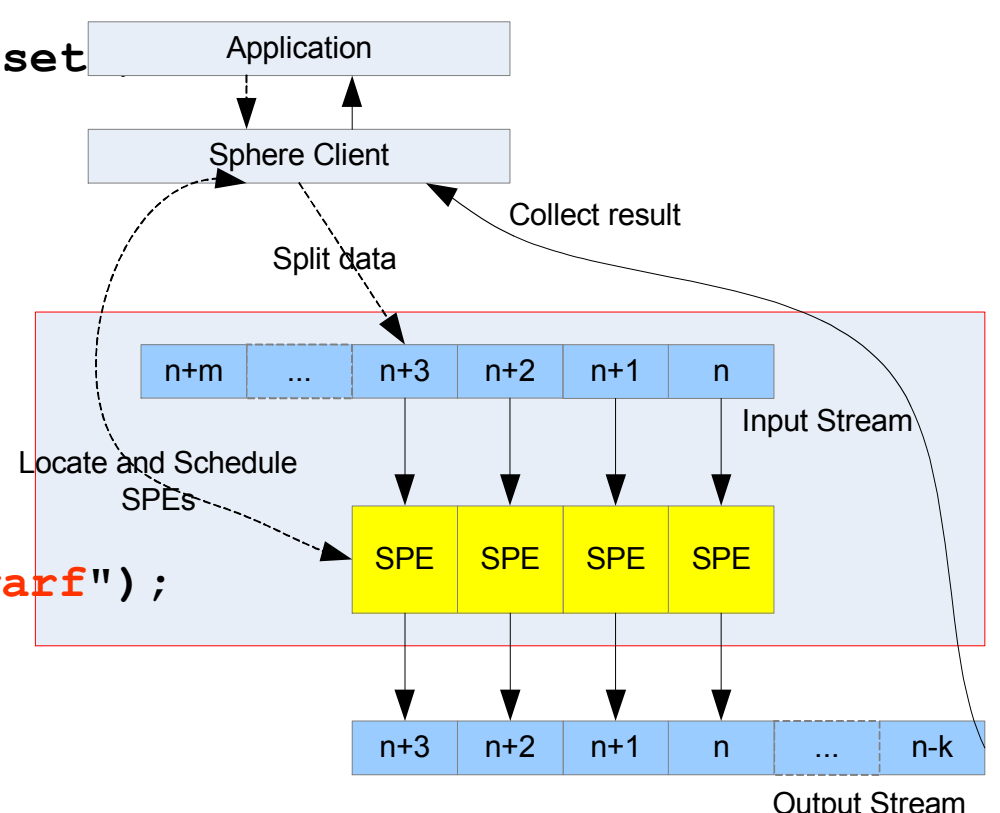
Sphere allows developers to write certain distributed data parallel applications with several simple APIs. User-define functions can be independently applied to each element in the dataset, while the result can be written to either the local disk or common destination files on other nodes. Data locality is a key factor for the performance of Sphere. In addition, sphere provides transparent load balance and fault tolerance.

```

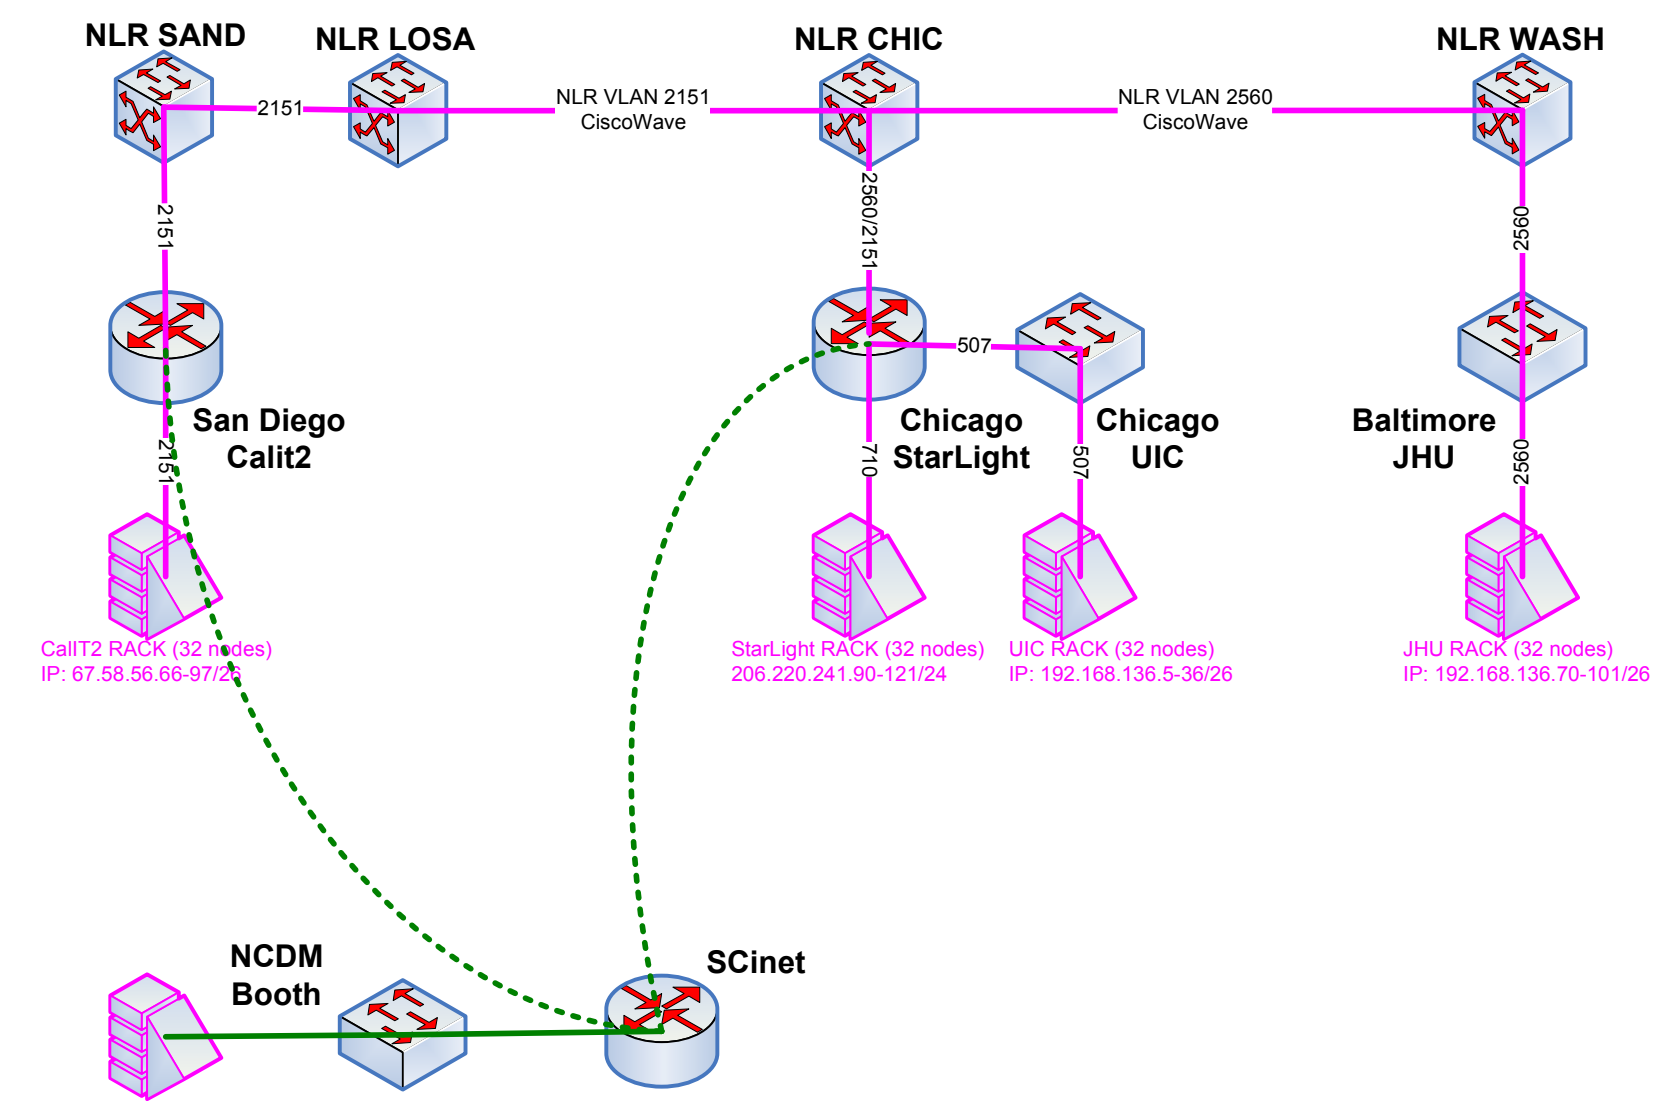
for each file F in (SDSS dataset)
  for each image I in F
    findBrownDwarf(I, ...);

SphereStream sdss;
sdss.init("sdss files");
SphereProcess myproc;
myproc->run(sdss, "findBrownDwarf");
myproc->read(result);

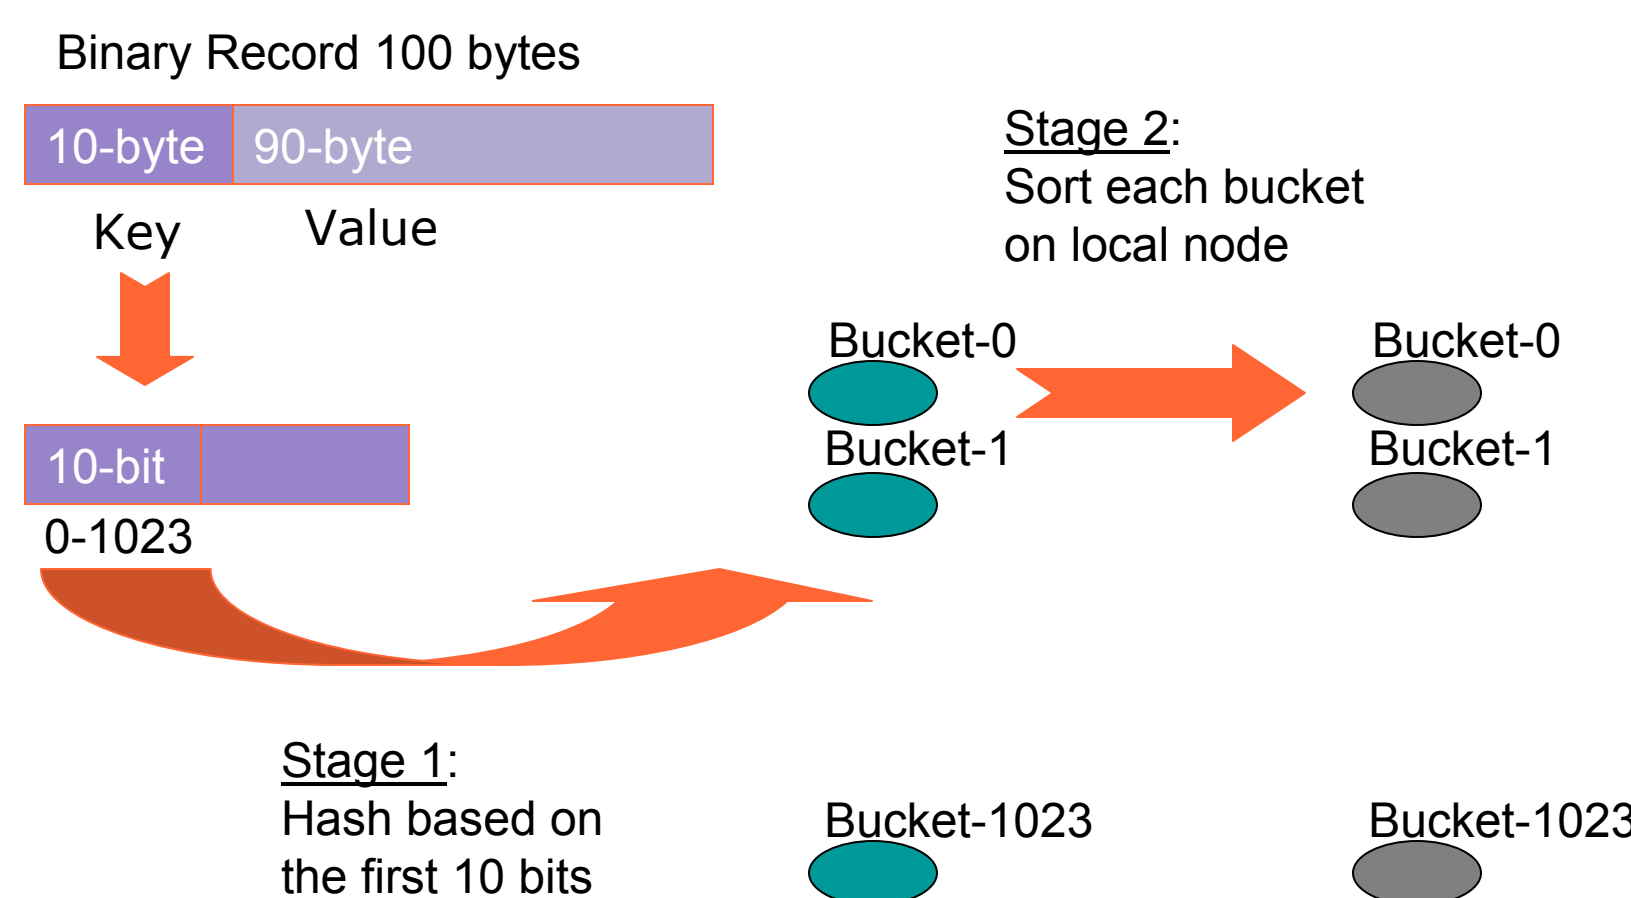
findBrownDwarf(char* image, int isize, char* result, int rsize);
    
```



The **Open Cloud Testbed** consists of 4 racks in 3 locations (Baltimore, Chicago/2, and La Jolla). Each rack has 32 nodes (dual AMD CPU, 12GB RAM, 1TB disk). The inter-rack bandwidth is 10Gb/s, provided by CiscoWave. The intra-rack bandwidth is 1Gb/s.

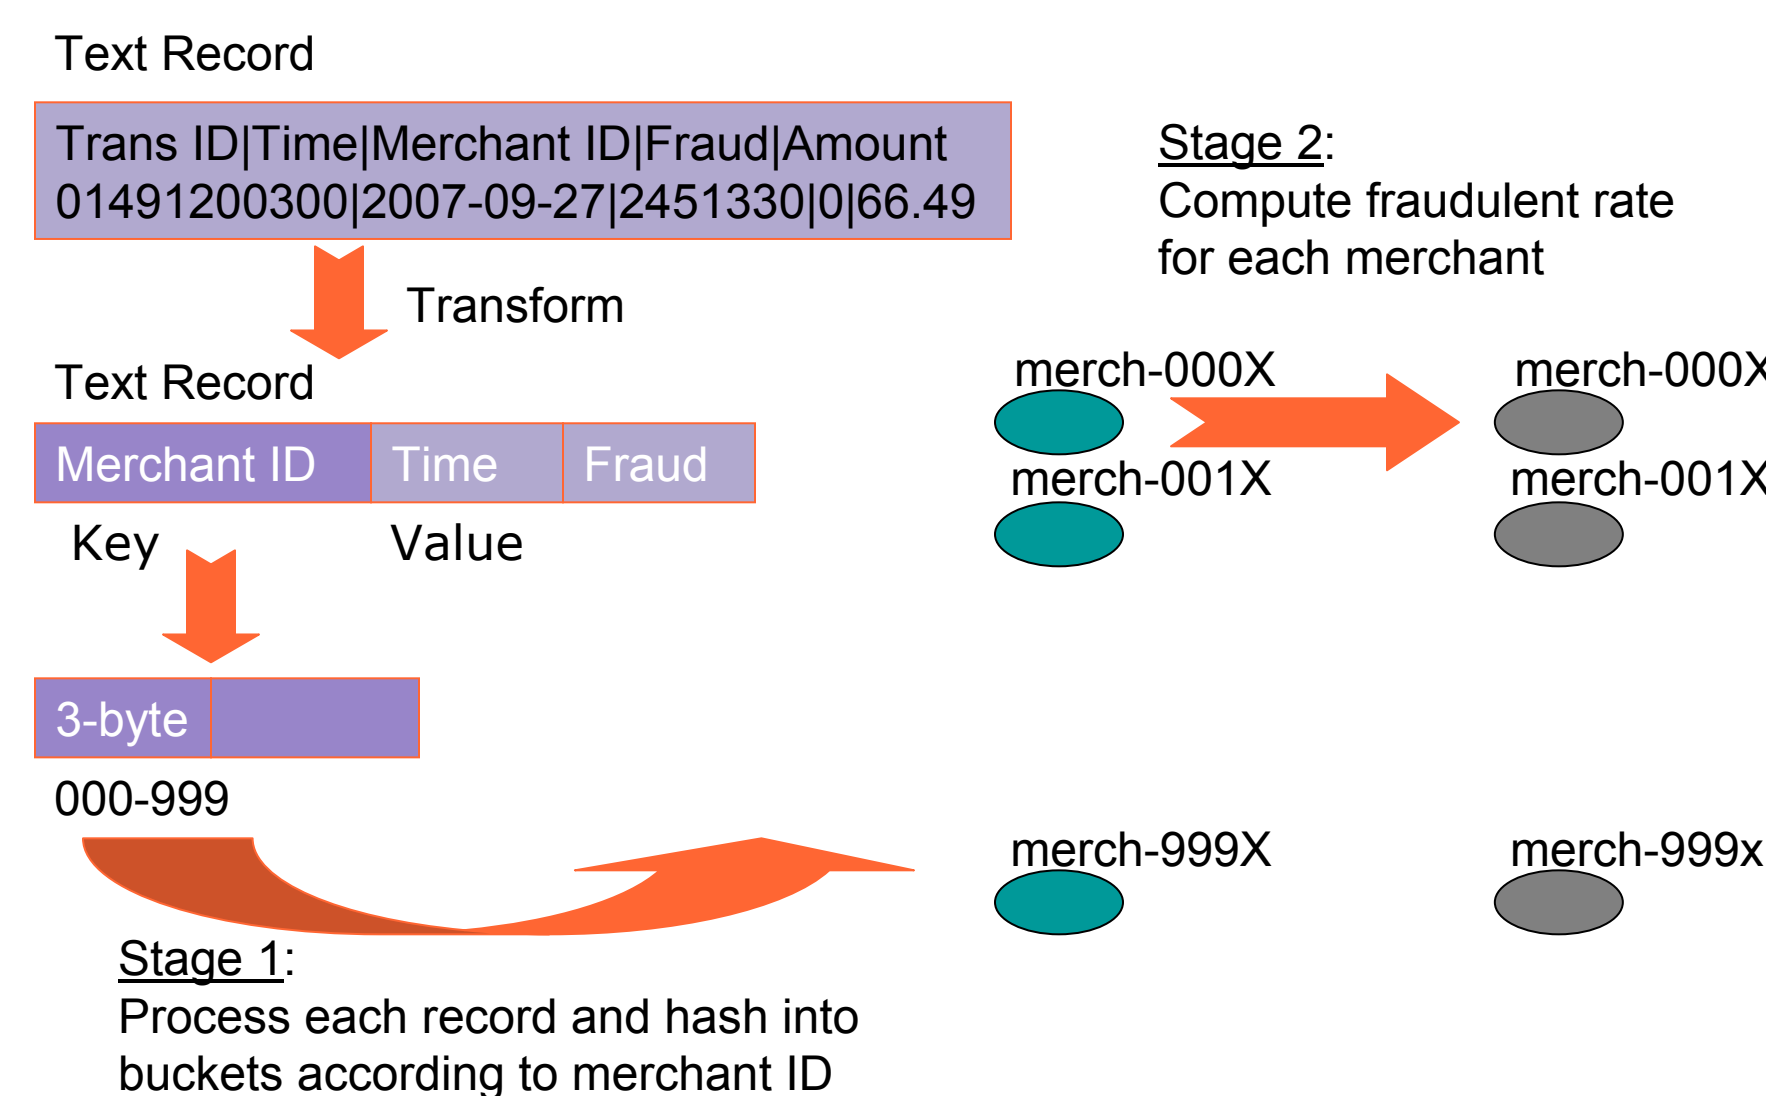


The **Terasort** benchmark sorts $N \times 10\text{GB}$ data, where N is the number of slave nodes in the system.



	Data Size	Sphere	Hadoop (3 replicas)	Hadoop (1 replica)
UIC	300GB	1265	2889	2252
UIC + StarLight	600GB	1361	2896	2617
UIC + StarLight + Calit2	900GB	1430	4341	3069
UIC + StarLight + Calit2 + JHU	1.2TB	1526	6675	3702

The **CreditStone** application computes the fraudulent rate for each merchant in the credit card transaction dataset.



Racks	JHU	JHU, SL	JHU, SL, Calit2	JHU, SL, Calit2, UIC
Number of Nodes	30	59	89	117
Size of Dataset (GB)	840	1652	2492	3276
Size of Dataset (rows)	15B	29.5B	44.5B	58.5B
Hadoop (min)	179	180	191	189
Sector with Index (min)	46	47	64	71
Sector w/o Index (min)	36	37	53	55