

SNIA COMPUTE + MEMORY
+ STORAGE SUMMIT

Architectures, Solutions, and Community
VIRTUAL EVENT, APRIL 11-12, 2023

Python with Computational Storage

Jayasankar OP, Associate Technical Director

Arun V Pillai, Staff Engineer

Sharanya P, Senior Software Engineer

Sangho Yi, Principal Engineer

Lakshmi Narayana, Senior Engineer

Samsung Semiconductor Inc




Agenda

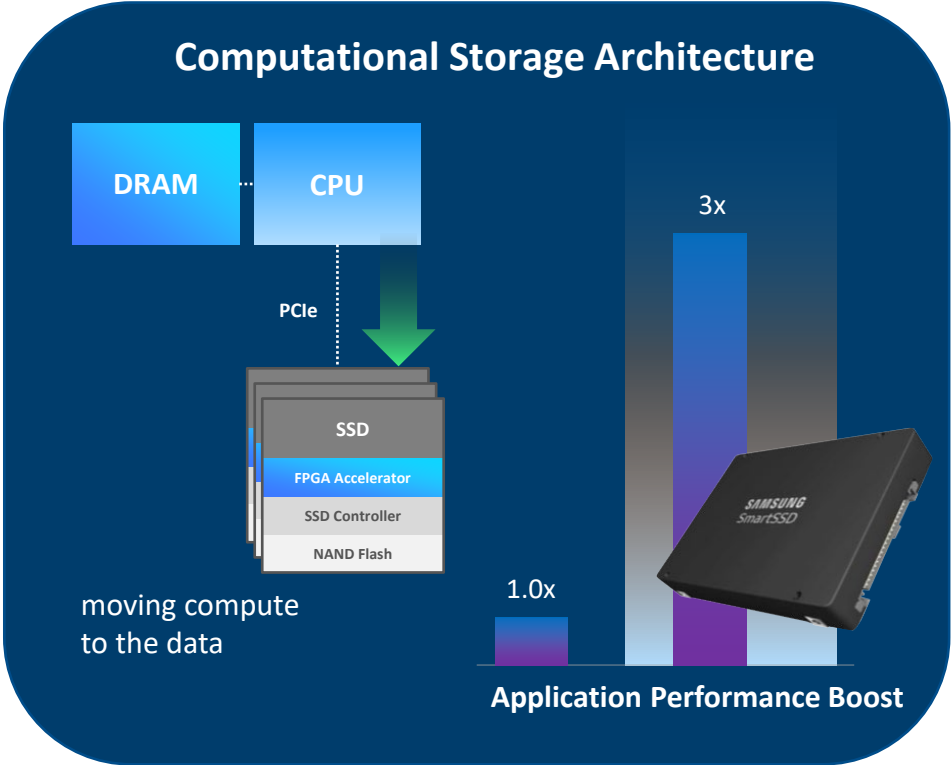
- Samsung SmartSSD[®]
- Software Stack for Computation Storage Drive (CSD)
- Python Interface to CSD
- Python CS Library
- A Python Based Unit Test Framework for CSx
- SmartSSD[®]: Geo Locator Using PostgreSQL
- What's Next ..?

Samsung SmartSSD[®]

- Samsung 2nd Gen SmartSSD based on NVMe standards in development (TP4091)

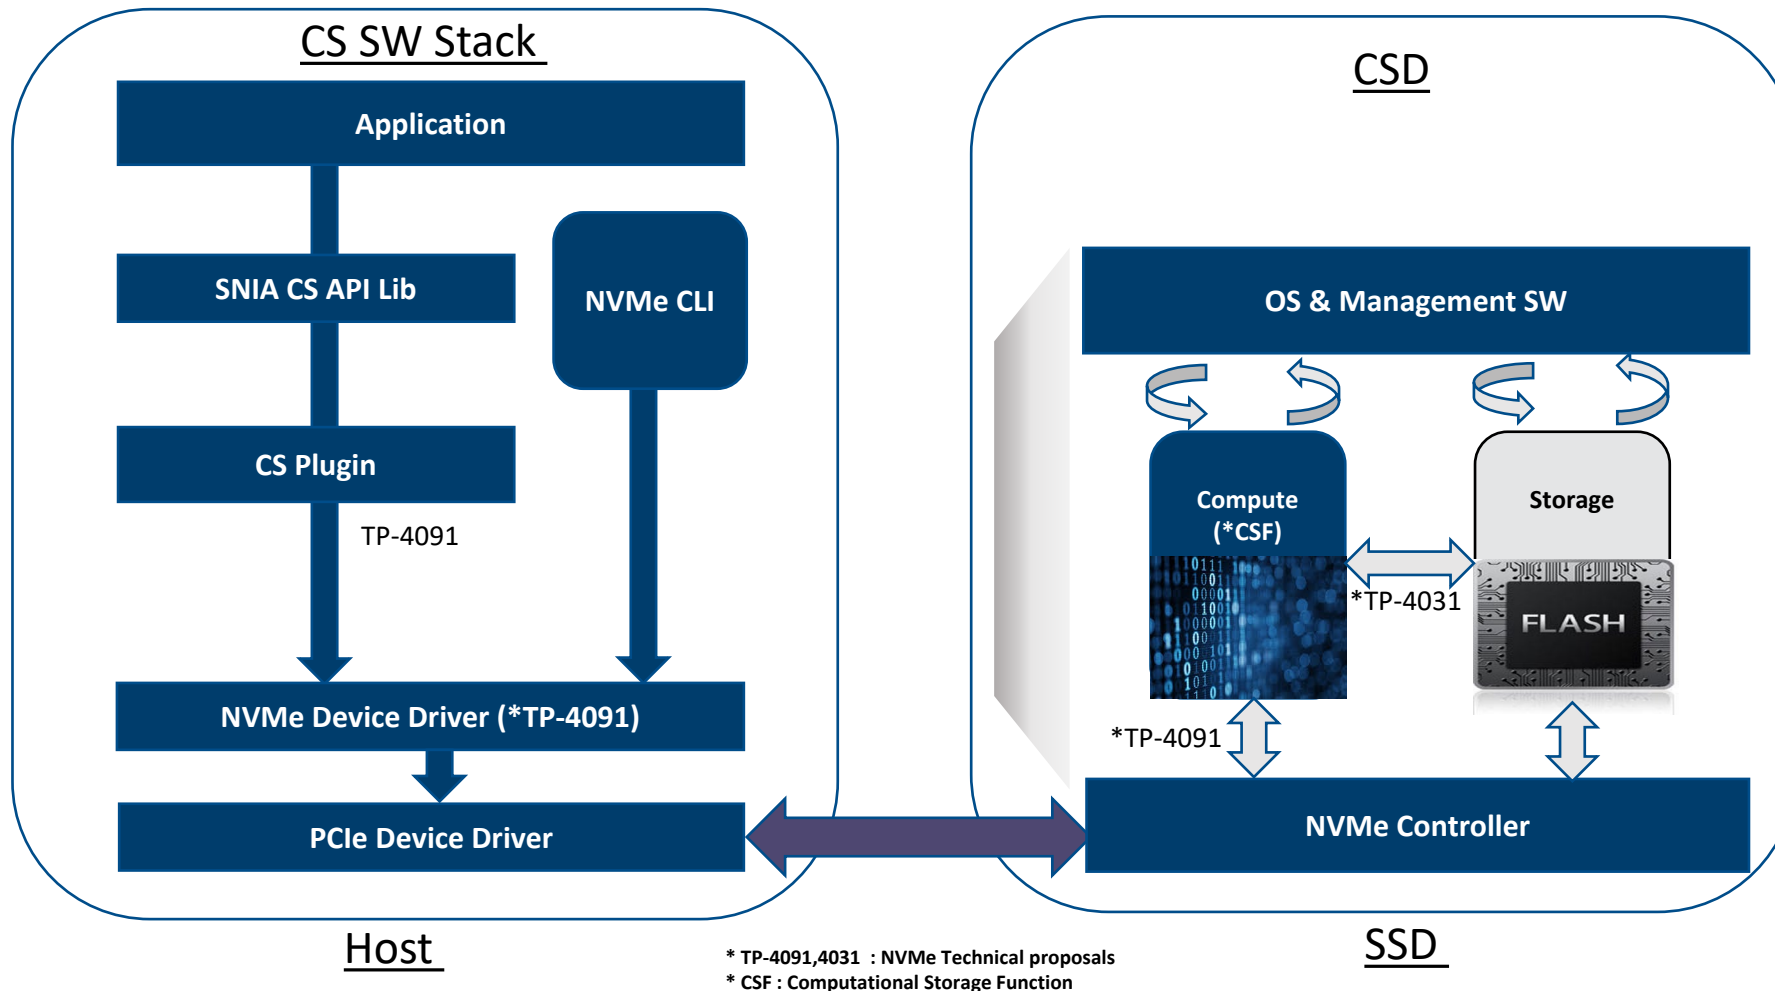


* SLM: Subsystem Local Memory



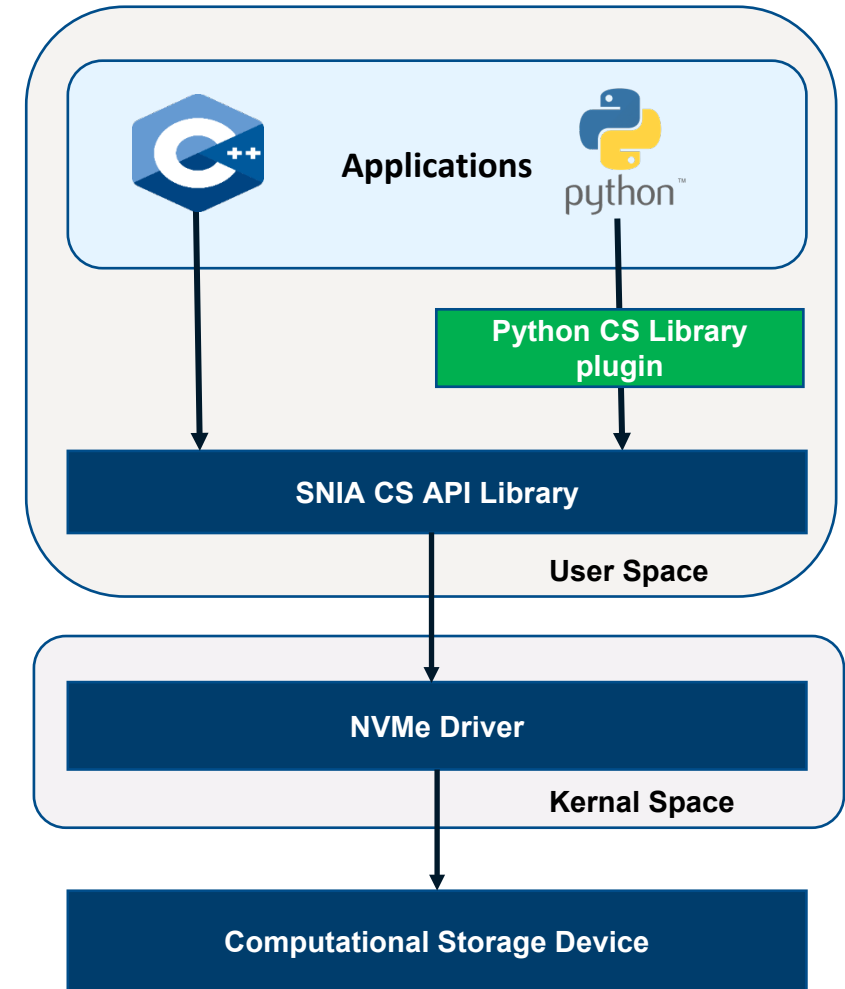
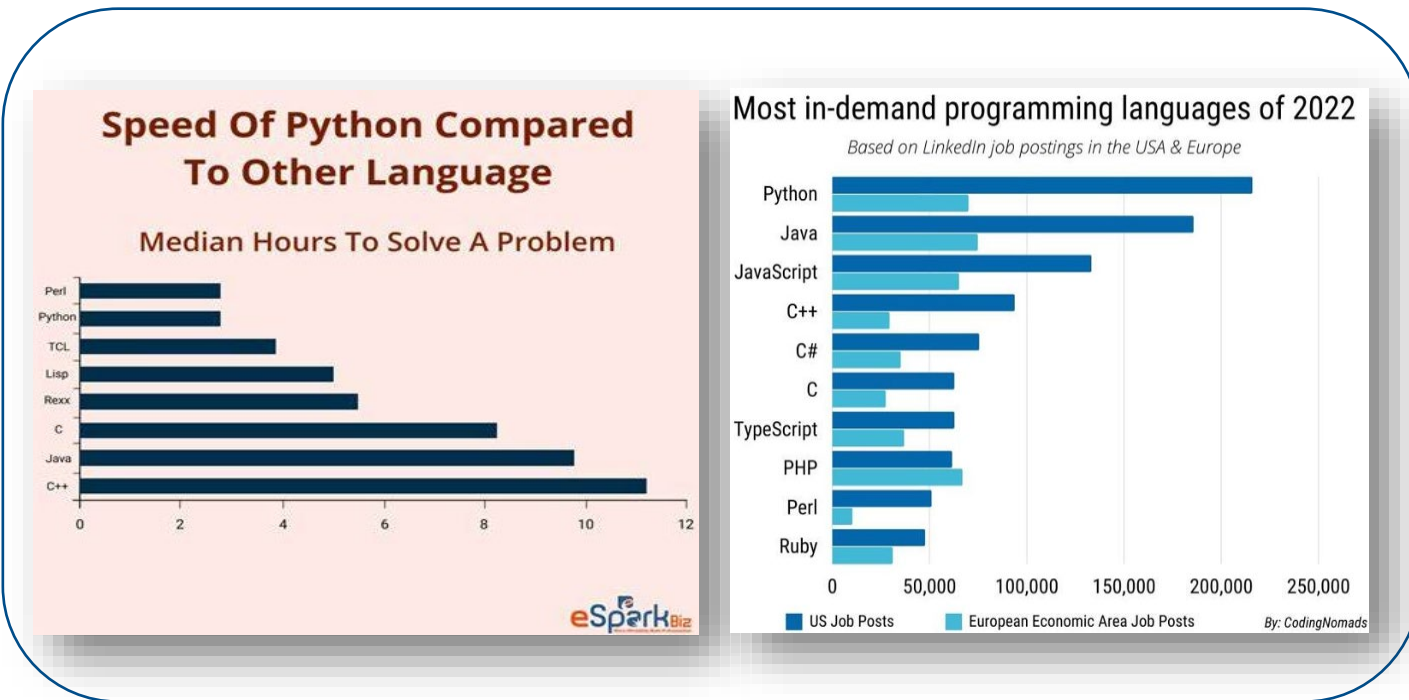
Software Stack for Computation Storage Drive (CSD)

- Computational Storage Drive capable of offloading host processing and achieve near-memory Computing



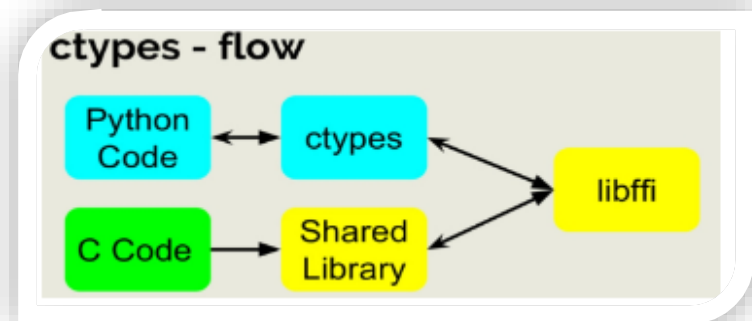
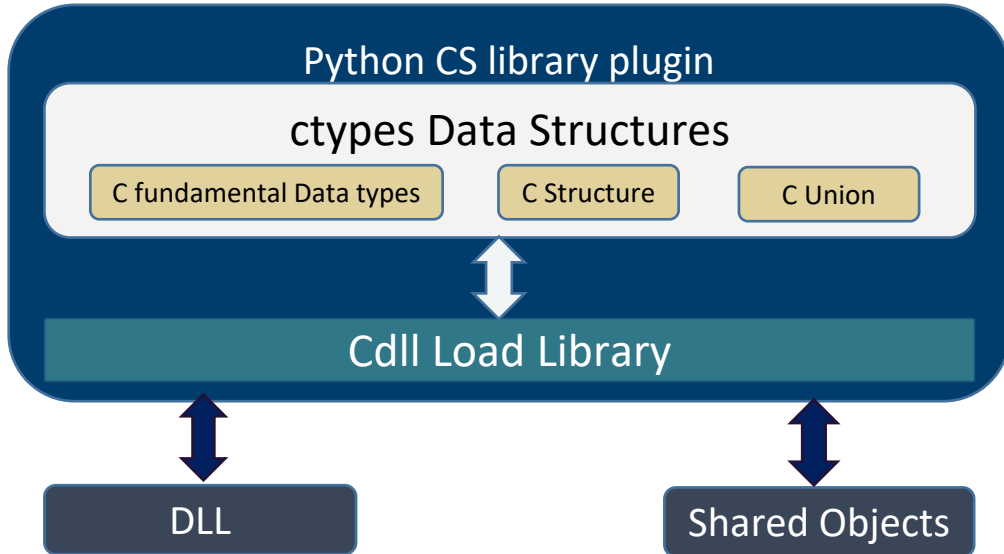
Python Interface with CSD

- Python leads on popularity
 - Top rated by demand/usage
 - Most desired by job listings
 - Ease of debug

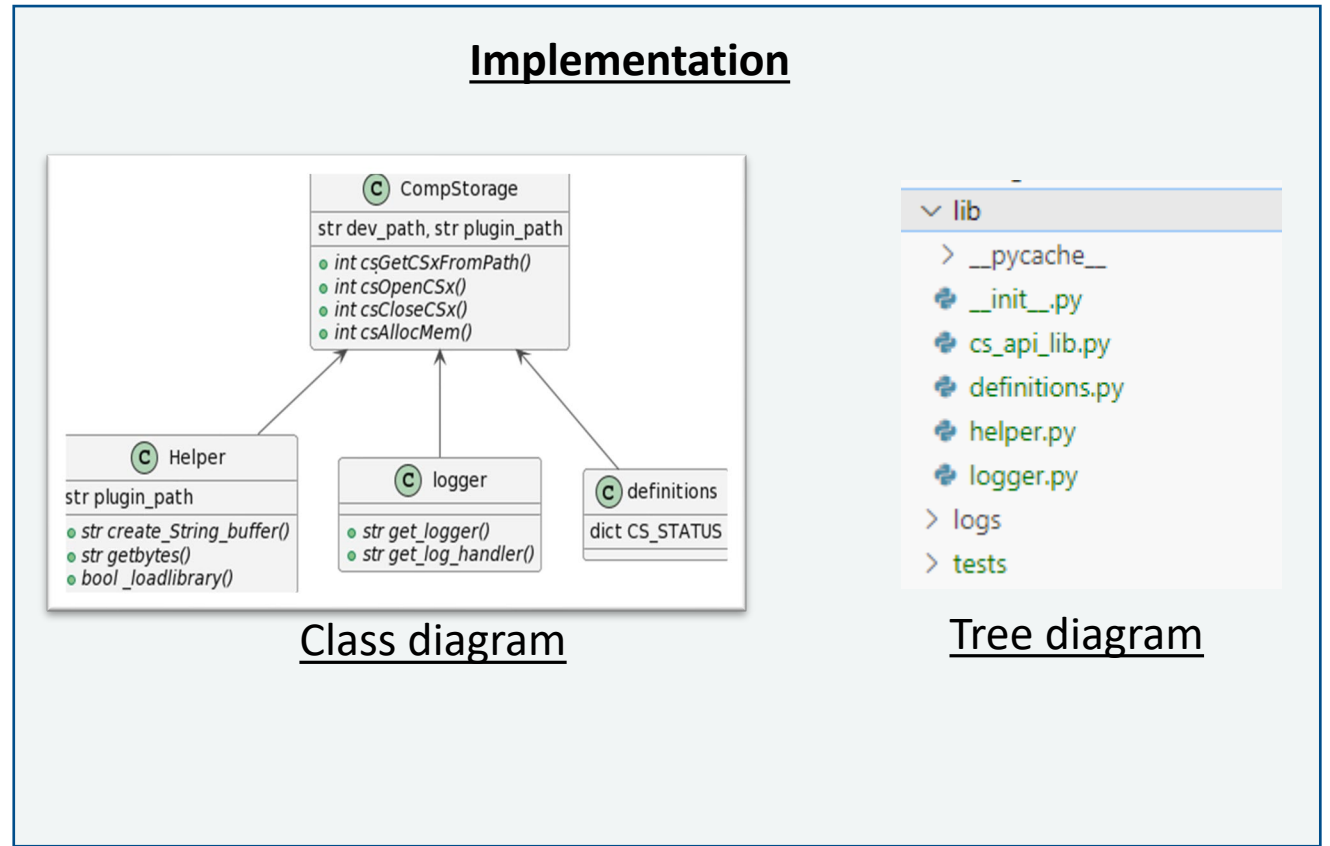


Python CS Library Plugin

- Python CS Library realized using Ctypes module

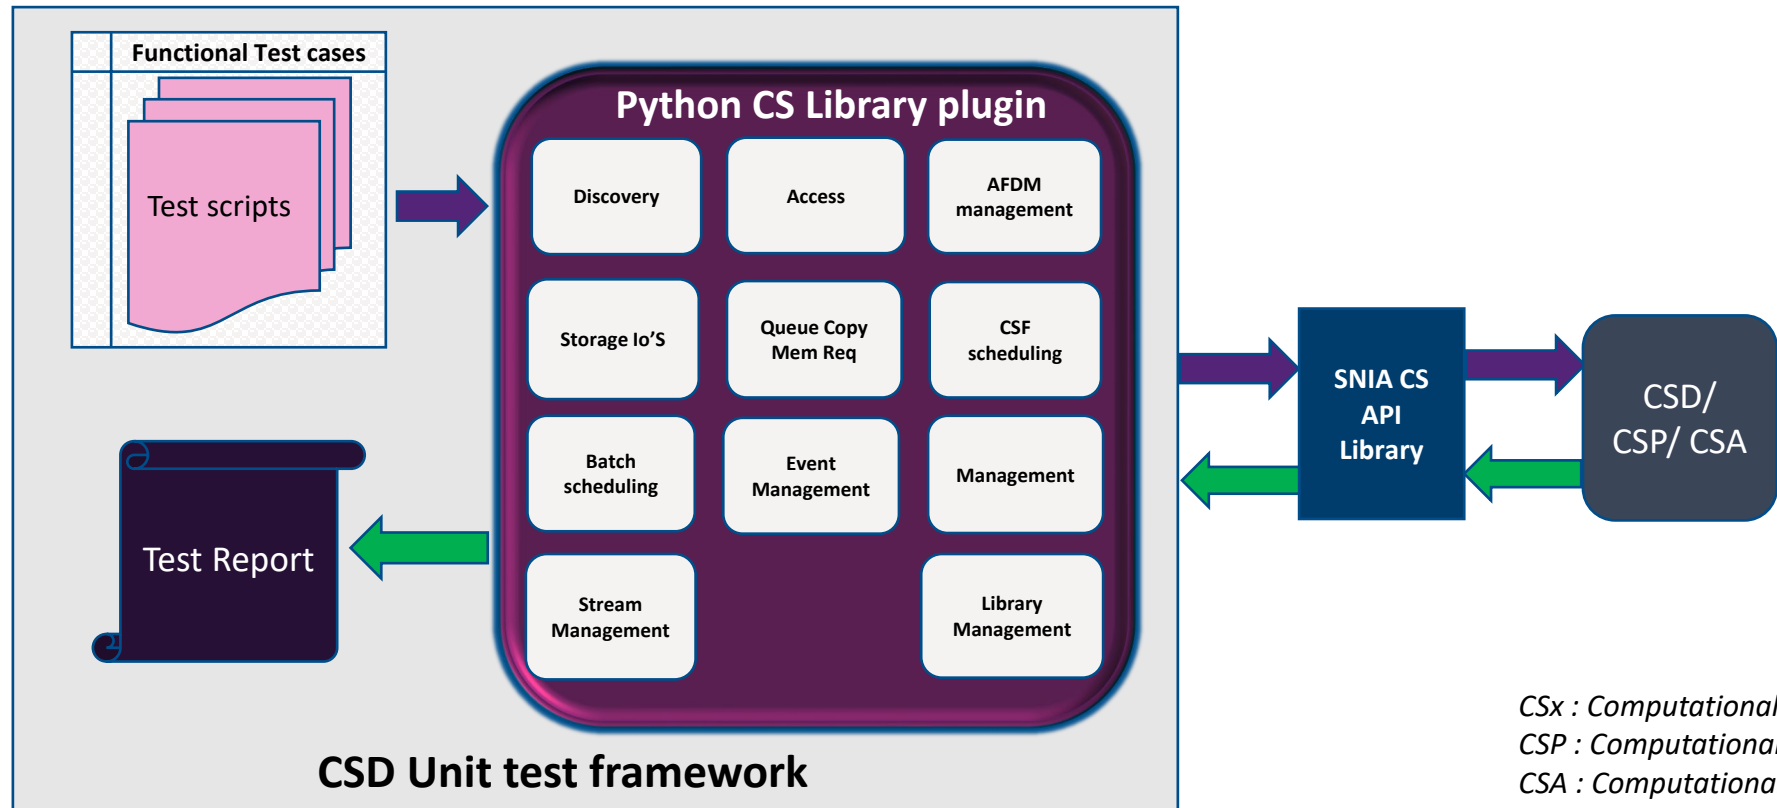


Implementation



Python Based Unit Test Framework for CS

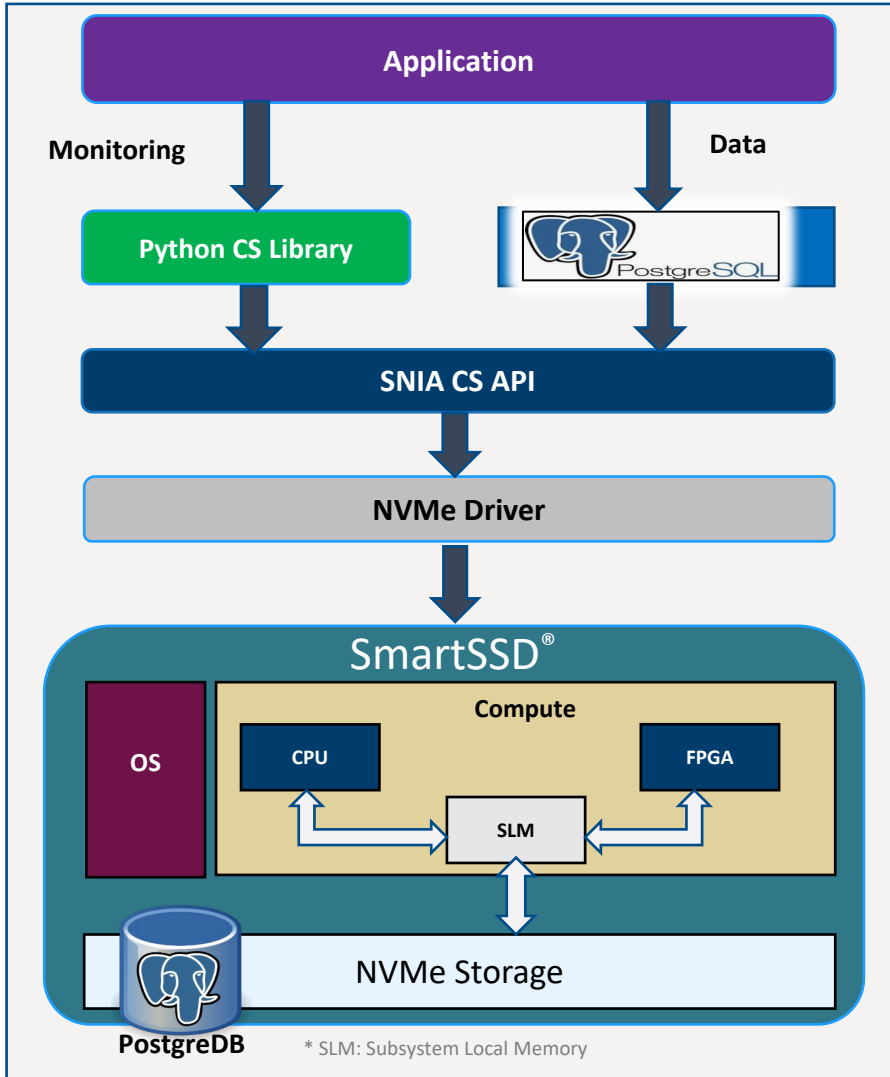
- Python based unit test suite to validate and verify functionality of CSx through SNIA CS API
- Implemented using python unit test framework



CSx : Computational Storage Device
CSP : Computational Storage Processor
CSA : Computational Storage Array

SmartSSD[®]: Geo Locator Using PostgreSQL

- PostGIS and spatial database extender enabled in PostgreSQL for geo location features



id	name	latitude	longitude	country	population	unique_id
48488	Tokyo	35.6839	139.7744	Japan	39105000	1392685764
48489	Jakarta	-6.2146	106.8451	Indonesia	35362000	1360771077
48490	Delhi	28.6667	77.2167	India	31870000	1356872604
48491	Manila	14.6	120.9833	Philippines	23971000	1608618140
48492	Sao Paulo	-23.5504	-46.6339	Brazil	22495000	1076532519
48493	Seoul	37.56	126.99	South Korea	22394000	1410836482
48494	Mumbai	19.0758	72.8775	India	22186000	1356226629
48495	Shanghai	31.1667	121.4667	China	22118000	1156073548
48496	Mexico City	19.4333	-99.1333	Mexico	22105000	1484247881
48497	Guangzhou	23.1288	113.259	China	21489000	1156237133
48498	Cairo	30.0444	31.2358	Egypt	19787000	1818253931
48499	Beijing	39.904	116.4075	China	19437000	1156228865
48500	New York	40.6943	-73.9249	United States	18713220	1840034016
48501	Kolkata	22.5727	88.3639	India	18698000	1356060520
48502	Moscow	55.7558	37.6178	Russia	17693000	1643318494
48503	Bangkok	13.75	100.5167	Thailand	17573000	1764068610
48504	Dhaka	23.7289	90.3944	Bangladesh	16839000	1050529279
48505	Buenos Aires	-34.5997	-58.3819	Argentina	16216000	1032717330
48506	Osaka	34.752	135.4582	Japan	15490000	1392419823
48507	Lagos	6.45	3.4	Nigeria	15487000	1566593751
48508	Istanbul	41.01	28.9603	Turkey	15311000	1792756324

DB table "citynames" : (41694 rows)

```

postgres=# select name, country, population, latitude, longitude
postgres=# from citydetails
postgres=# where ((latitude between -5 and 5) and population > 2000000;
name | country | population | latitude | longitude
-----+-----+-----+-----+-----
Kinshasa | Congo (Kinshasa) | 15056000 | -4.3317 | 15.3139
Kuala Lumpur | Malaysia | 8639000 | 3.1478 | 101.6953
Bogota | Colombia | 7743955 | 4.6126 | -74.0705
Nairobi | Kenya | 5545000 | -1.2864 | 36.8172
Singapore | Singapore | 5271000 | 1.3 | 103.8
Timbio | Colombia | 4444444 | 2.3528 | -76.6819
Guayaquil | Ecuador | 2723665 | -2.19 | -79.8875
Cali | Colombia | 2471474 | 3.44 | -76.5197
Fortaleza | Brazil | 2452185 | -3.7275 | -38.5275
Yaounde | Cameroon | 2440462 | 3.8578 | 11.5181
Douala | Cameroon | 2446945 | 4.05 | 9.7
Mogadishu | Somalia | 2120000 | 2.0408 | 45.3425
Medan | Indonesia | 2109330 | 3.6667 | 98.6667
Quito | Ecuador | 2011388 | -0.22 | -78.5125
(14 rows)
Time: 9.738 ms
    
```

```

postgres=# select name, country,
postgres=# (
postgres=# point(-118.4068, 34.1139)<=>point(local.longitude, local.latitude)
postgres=# ) * 1609.34::int as point_point_distance_in_meters
postgres=# from citydetails,
postgres=# lateral (
postgres=# select id,latitude, longitude from citydetails
postgres=# where ((latitude between 33.1139 and 35.1139)
postgres=# and (longitude between -119.4068 and -118.84297))
postgres=# ) as local
postgres=# where citydetails.id = local.id;
name | country | point_point_distance_in_meters
-----+-----+-----
Oxnard | United States | 71855.61647291823
Thousand Oaks | United States | 43972.17997046268
San Buenaventura | United States | 77888.34782426493
Santa Paula | United States | 66442.11945303493
Camarillo | United States | 58789.97434961202
Moorpark | United States | 47255.503117375905
Port Hueneme | United States | 73509.76226287287
Fillmore | United States | 56666.34332526211
Ojai | United States | 85676.85000993097
Mira Monte | United States | 87954.80720469833
El Rio | United States | 70504.56735180593
(11 rows)
Time: 9.507 ms
    
```

Find cities near earth equator having population more than 2000000

Calculate point-point distance between Los Angeles and near cities in meters

Sample data source : <https://simplemaps.com/data/world-cities>

What's Next ?..

- Python CS Library performance tuning
- Unit test case coverage for all CSx workflow



COMPUTE + MEMORY + STORAGE SUMMIT

Architectures, Solutions, and Community
VIRTUAL EVENT, APRIL 11-12, 2023



Please take a moment to rate this session.

Your feedback is important to us.