



OpenMP 4.5 Validation and Verification Suite for Device offload

Sunita Chandrasekaran Jose Monsalve Diaz

Swaroop Pophale

Oscar Hernandez

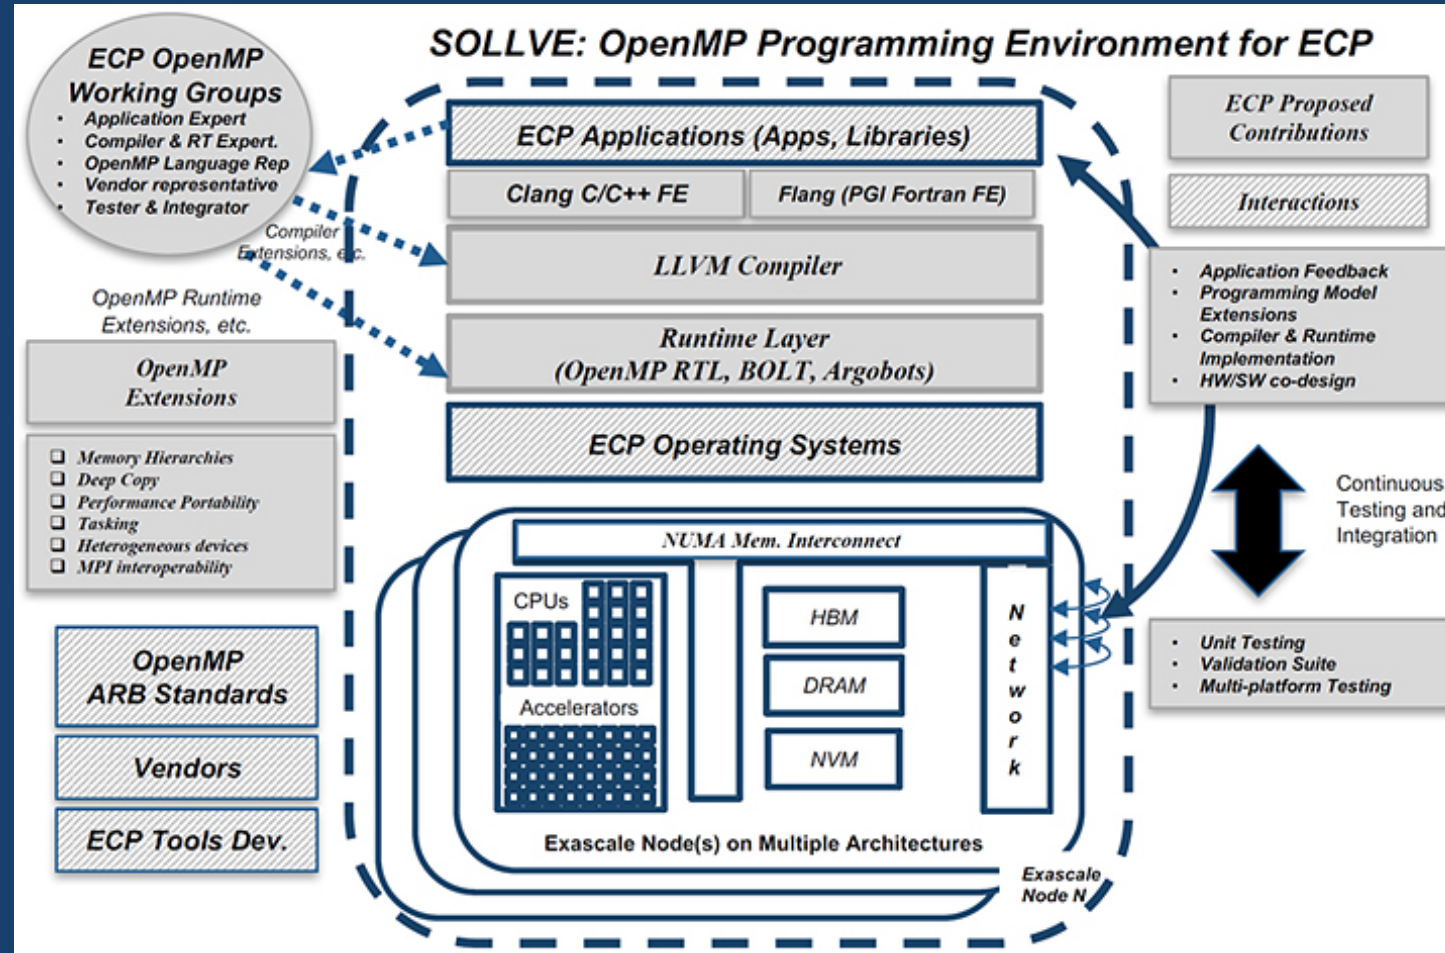
David E. Bernholdt

Outline


- Introduction
 - Problem Statement and Contributions
 - Related work
- OpenMP 4.5 offloading
- Methodology
 - Test design
 - Infrastructure design
 - Results logs and reports
- Results

Introduction

The ECP SOLLVE Project



Why Accelerator devices?



- Power wall limited single core performance and big cores
- Simpler cores and considerably larger core count
- But there are still programmability challenges...

Problem Statement

As the OpenMP specification continues to grows, **what are the methodologies and tools** to measure the level of compliance of a compiler implementation and supported system with respect to the OpenMP specification document?

Motivation

Programming model specifications is a


“Legal” document that binds **compiler implementation** and **users**

- Compiler developers use the specification to define the compiler behavior. To claim support the specs must be respected.
- Users do not need to learn all the implementation specific aspects to use the programming model as long as they know the specs

What about the system where the user is running on?



Motivation



Why not ____ ?

○ Examples as tests?

- Lack of coverage
- Different purpose

○ Automatic testing?

- Language interpretation of the Specs (Human required)


○ Use vendor tests?

- Vendor specific methodology
- Biased interpretation

OpenMP 4.5 Offloading model

OpenMP 4.5 Machine model

HOST



Host Memory

Interconnect

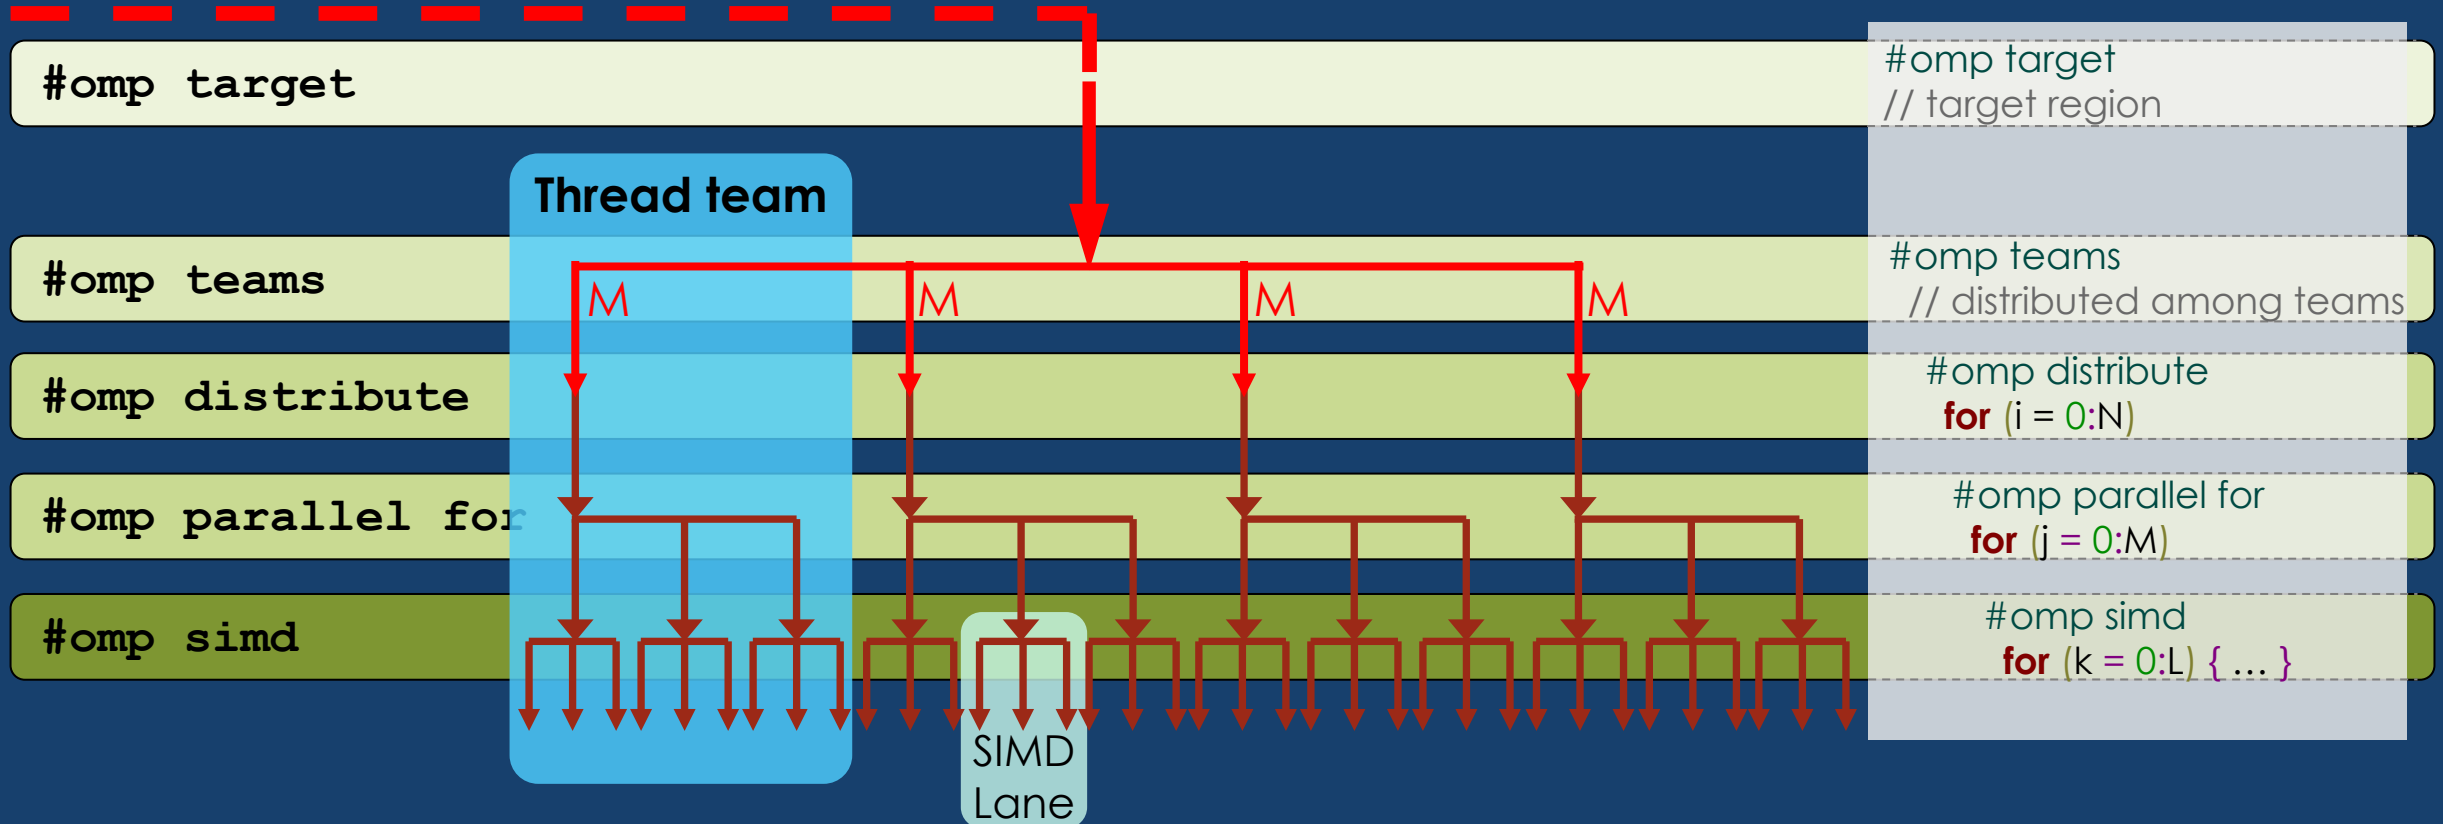
Device 1
Memory

Device N
Memory

DEVICES 11

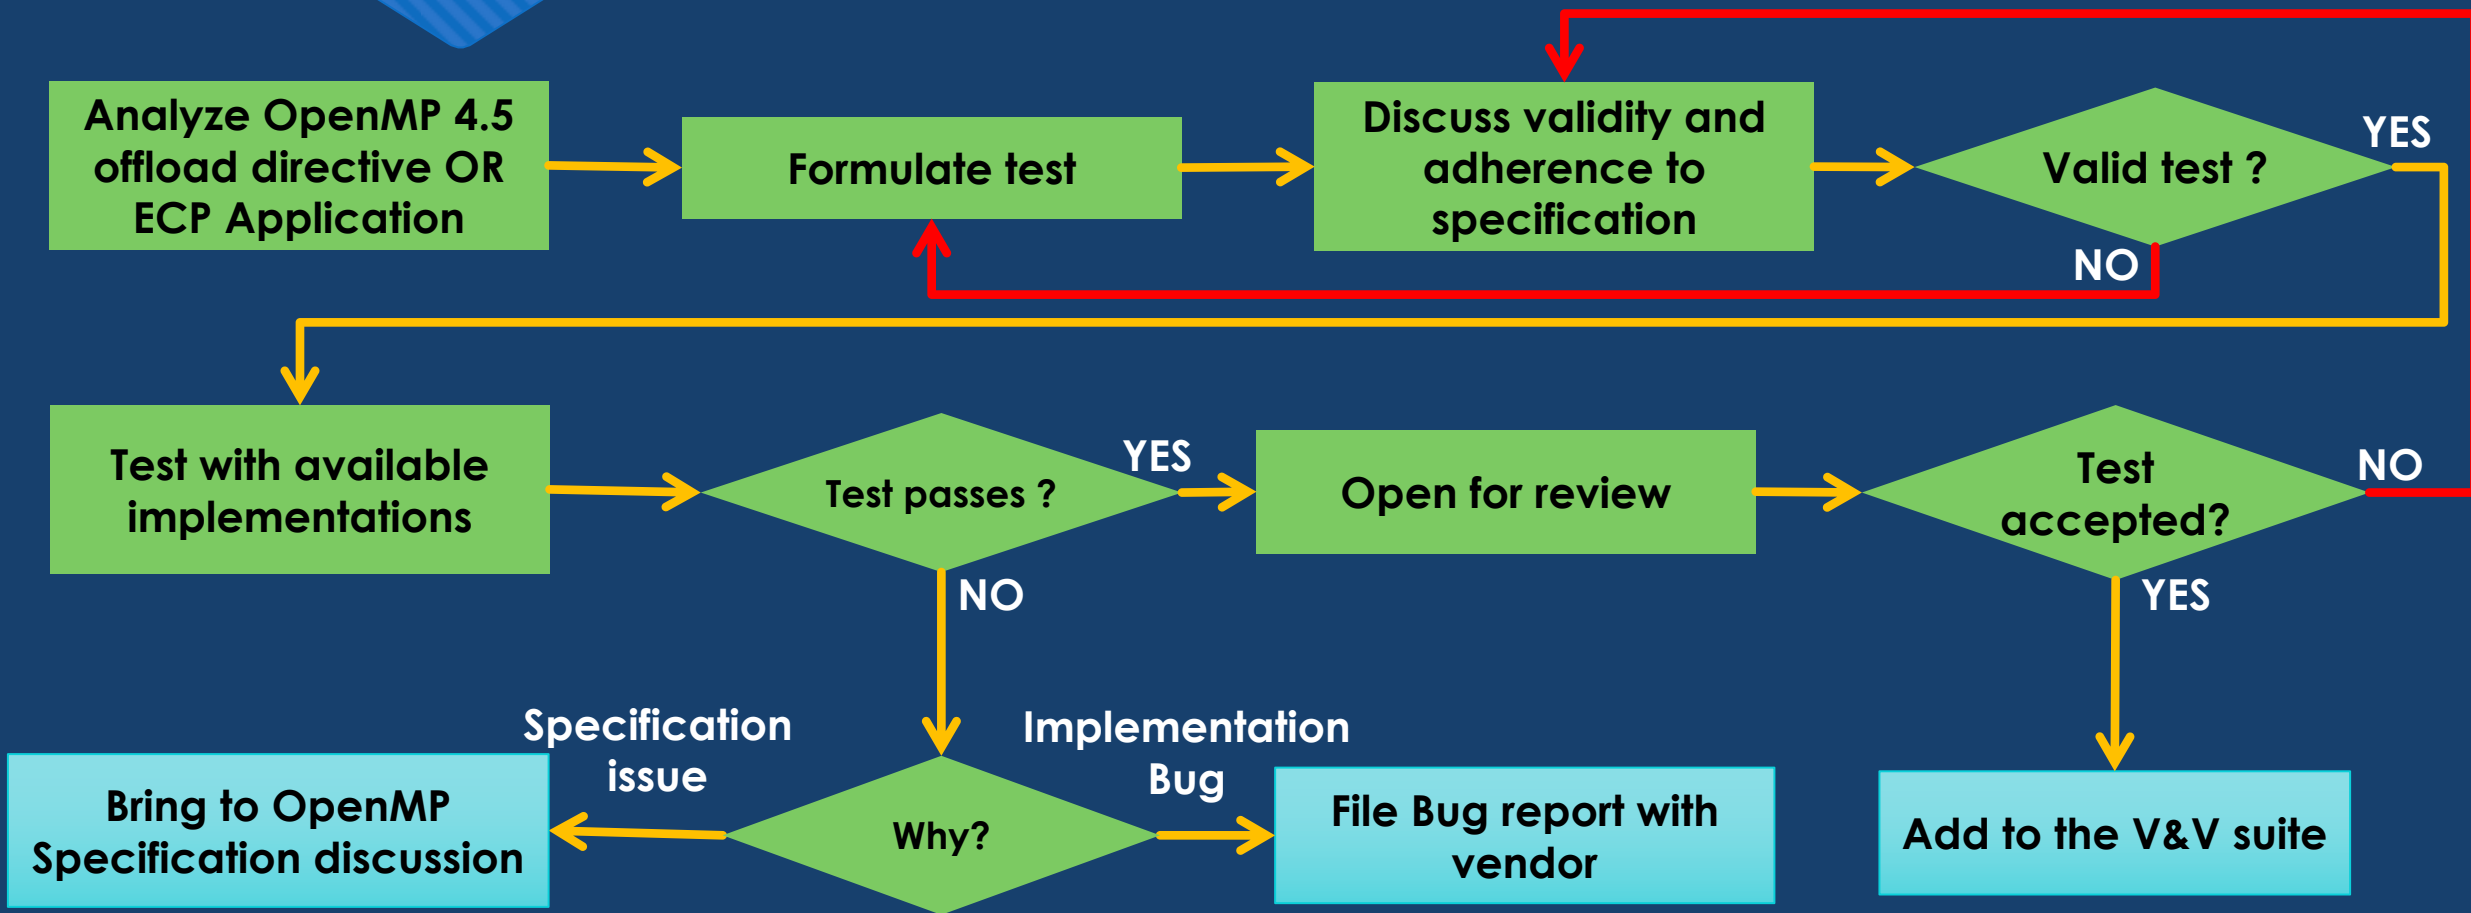
OpenMP 4.5 Execution Model

HOST to TARGET OFFLOADING



Tests suite methodology

Methodology



Test design

Data mapping on multiple devices

```
1  int num_dev = omp_get_num_devices();
2  int h_matrix [N*num_dev];
3
4  for (int dev = 0; dev < num_dev; ++dev) {
5  #pragma omp target data map(from: h_matrix[dev*N:N]) device(dev)
6  {
7  #pragma omp target map(from: h_matrix[dev*N:N]) device(dev)
8  {
9      for (int i = 0; i < N; ++i)
10         h_matrix[dev*N + i] = dev;
11 } // end target
12 } // end target data
13 } // end for loop
14
15 // checking results
16 for (int dev = 0; dev < num_dev; ++dev) {
17     for (int i = dev*N ; i < (dev+1)*N; ++i)
18         OMPVV_TEST(errors, dev != h_matrix[i]);
19 }
```

Get number of available devices

Map data from selected device to host

Create target region on selected device

Computation offloaded to selected device

Compare results with expected

Test design

Target teams distribute firstprivate (segment)

```
1 int a[N] = 1; b[N] = x; c[N] = 2*x; d[N] = 0;
```

Variable to privatize

```
2 int privatized = 1;
```

```
3 #pragma omp target teams distribute firstprivate(privatized) \
```

Privatize and initialize value

```
4 map(from: d[:]) map(to: a[:], b[:], c[:])
```

```
5 for (int x = 0; x < N; ++x) {
```

```
6   for (int y = 0; y < a[x] + b[x]; ++y) {
```

```
7     privatized++;
```

Possible data race if privatization fails

```
8   }
```

```
9   d[x] = c[x] * privatized; privatized = 0;
```

d[] serves as probe to check the value

```
10 }
```

```
11
```

```
12 for (int x = 0; x < N; ++x){
```

```
13   OMPVV_TEST_AND_SET_VERBOSE(errors, d[x] != (2 + x) * 2 * x);
```

Check the expected value of d[]

```
14 }
```

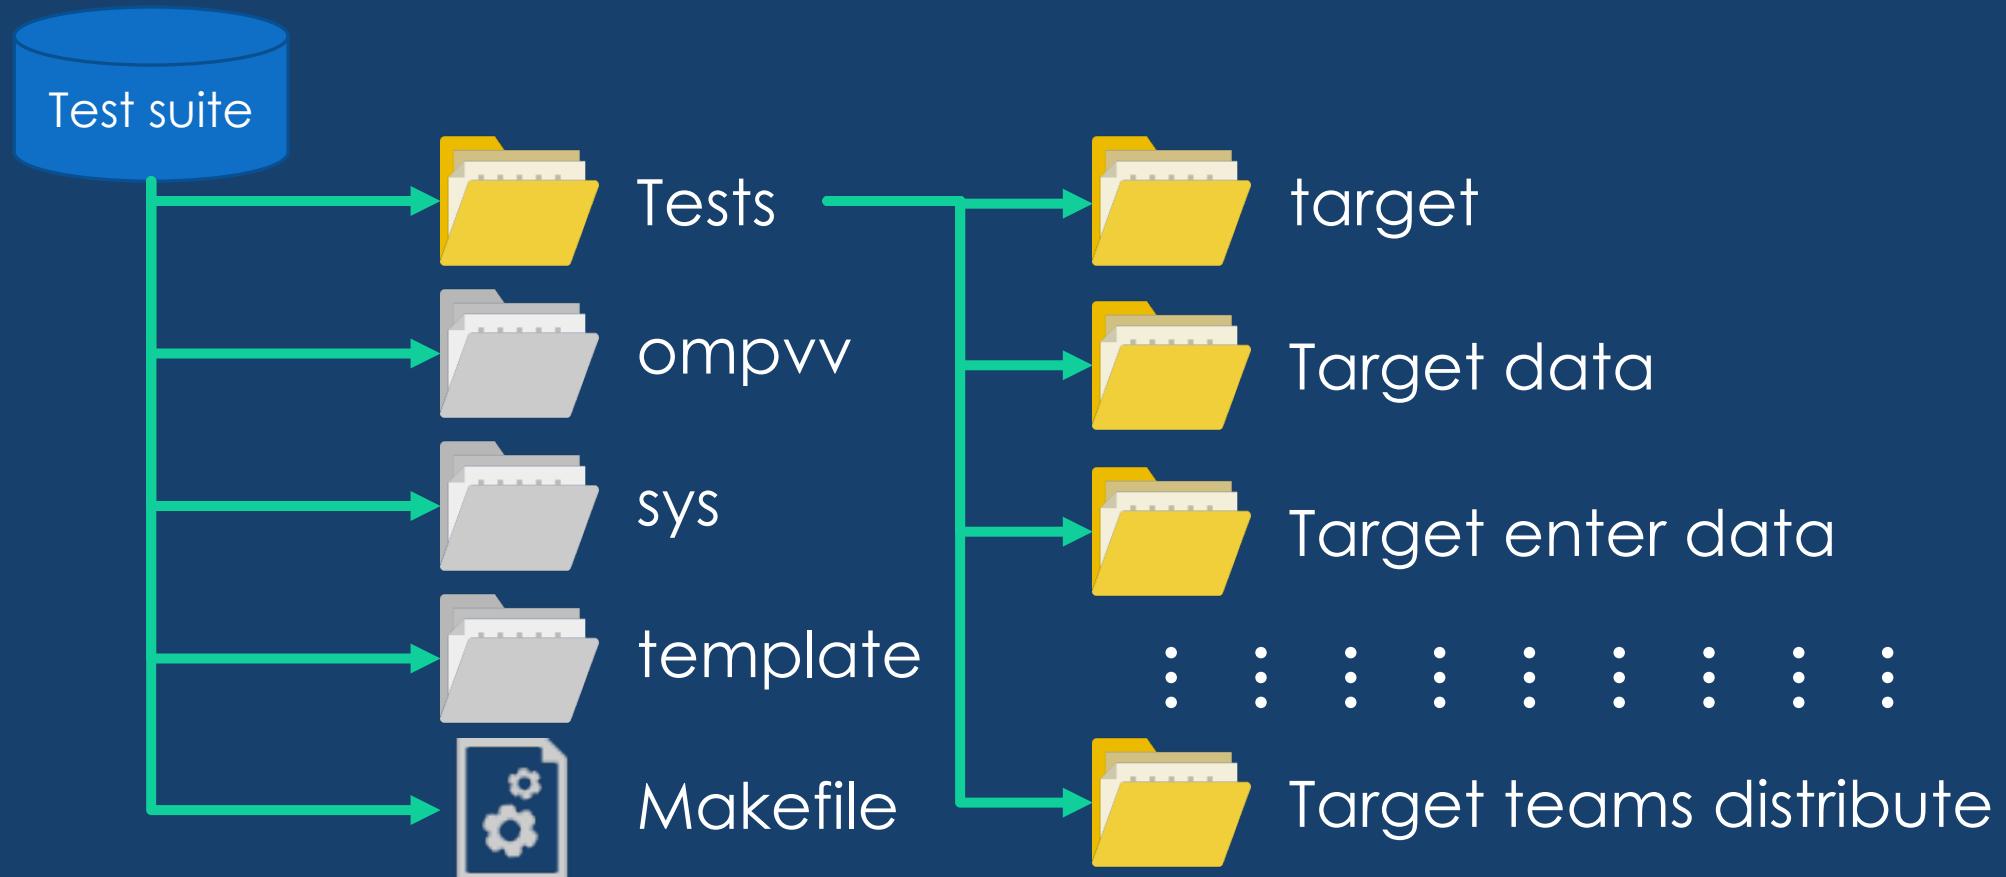
Infrastructure design

Our infrastructure is based on a **Makefile + scripts**

Design parameters:

- Portability across multiple **compilers** and **systems** and easy to use
 - Support for different compiler options
 - Support for Lua-like Modules and batch schedulers (Usually used in HPC clusters)
- Fast test addition and modification
- Divided compilation and execution phases
- Subset of tests selection for partial execution

Folder structure



Makefile Options

Option	Description
VERBOSE=1	Make output verbosity
VERBOSE_TESTS=1	Test output verbosity
LOG=1	Enable test output logging
LOG_ALL=1	Enables Make output and test output logging
SYSTEM=sys_name	Include system definition file
ADD_BATCH_SCHED=1	Add batch schedule line before running a test
MODULE_LOAD=1	Add <i>module load</i> commands before every command
SOURCES=file.(c,cpp,F90)	Select a subset of tests

Makefile Simple Example



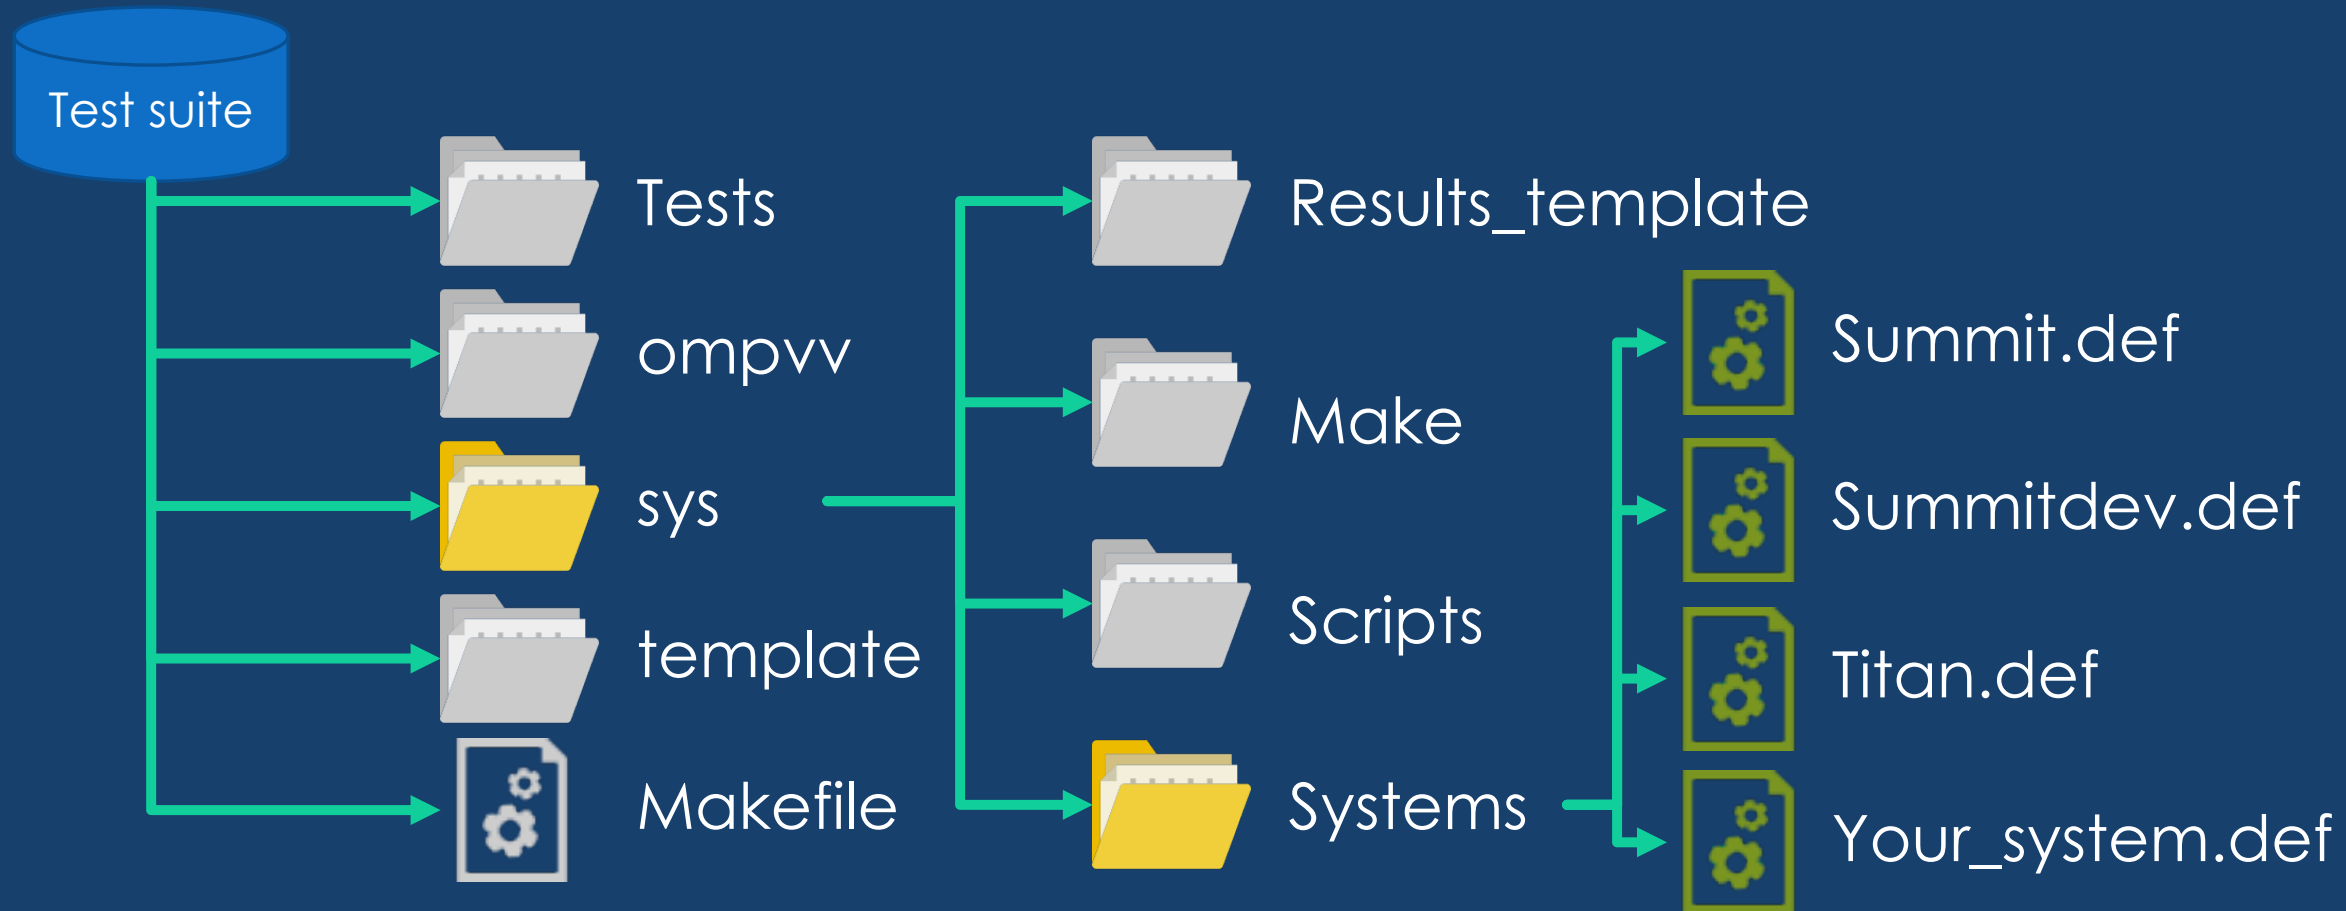
Makefile Rules

Option	Description
all	Compile and run
compile	Compile only
run	Run previously compiled tests
clean	Clean previously compiled tests
compilers	List available compilers (SYSTEM option dependent)
report_json	Parse raw logs and produce results.json file
report_html	Use JSON file and template to produce html visualization of results

System support

- Specify system specific support for:
 - Linux-line Module loading for each compiler
 - Batch scheduling execution lines (e.g. jsrun and aprun)
 - Get compiler version (command)
 - Specify compiler specific flags (CFLAGS, CXXFLAGS)

Systems folder



Results

Log files

- 3 Different Log formats:
 - RAW Format:
 - Format specific output log.
 - Single file per test,
 - Single section per action delimited by header and footer with system and runtime information
 - JSON FILE:
 - Translation process from parsing the raw format
 - HTML (visual):
 - Template that uses the JSON file to load results into a visualization table

Experimental setup

System	Model	Processors	Cores/node	Threads/node	Memory	Accelerator	Compilers
Titan	Cray XK7	AMD Opteron 6274	16	16	32 GB	1 NVIDIA K20X	CCE 8.7.2
Summitdev	IBM S822LC	2x Power8	20	160	256 GB	4 NVIDIA P100	GCC 7.1.1 Clang 3.8.0 XLC 13.1.6
Summit	IBM AC922	2x Power9	42	168	512 GB	6 NVIDIA V100	Clang 3.8.0 XLC 13.1.7

Current results snapshots

Visit our website for more

OpenMP Features	Summitdev					Summit **		Titan
	GCC 7.1.1	gfortran 7.1.1	Clang CORAL 3.8.0	XLC 13.1.6	XLF 15.1.7	Clang CORAL 3.8.0	XLC 13.1.7	CCE 8.7.2
target	14/15	13/13	15/15	14/15	12/14	12/14	11/14	13/14
target data	5/5	10/10	6/6	6/6	8/10	6/6	6/6	3/6
target enter/exit data	6/6	4/4	6/6	6/6	5/5	6/7	6/7	5/7
target enter data	7/7	6/6	7/7	7/7	7/7	6/7	6/7	5/7
target update	5/5	-	5/5	5/5		5/5	4/5	4/5
target teams distribute **	12/18	2/8	10/18	12/18	2/8	-	-	9/11
target teams distribute parallel for **	13/14	-	13/14	11/14	-	-	-	10/14

Example Implementation discrepancy

Specifications (Section 2.15.5 - page 216):

- 4 • If a **defaultmap(tofrom:scalar)** clause is not present then a scalar variable is not
5 mapped, but instead has an implicit data-sharing attribute of firstprivate (see Section 2.15.1.1 on
6 page 179).

Test:

```
enum { VAL1 = 1, VAL2, VAL3, VAL4} scalar_enum = VAL1
```

```
#pragma omp target
```

```
{  
    scalar_enum = VAL4;  
}
```

```
OMPVV_TEST_AND_SET_VERBOSE(errors, scalar_enum != VAL1);
```

Failed condition



Visit our website

<https://crpl.cis.udel.edu/ompvvsolve/>



Contact information:

Jose Monsalve (josem@udel.edu)

Swaroop Pophale (pophale@ornl.gov)

Kyle Friedline (utimatu@udel.edu)

Oscar Hernandez (oscar@ornl.gov)

Sunita Chandrasekaran (schandra@udel.edu)



Work supported by the **U.S. Department of Energy**, Office of Science, the **Exascale Computing Project (17-SC-20-SC)**, a collaborative effort of the **U.S. Department of Energy Office of Science** and the **National Nuclear Security Administration** under contract number **DE-AC05-00OR22725**.

Copyright

Gavel taken from(Chris Potter - Modifier: Ibrahim.ID):

https://commons.wikimedia.org/wiki/File:3D_png_Judges_Gavel.png

Logos under CC0:

- <https://pixabay.com/en/papers-stack-heap-documents-576385/>
- <https://pixabay.com/en/gear-wheel-gearwheel-gear-cogs-310906/>
- <https://pixabay.com/en/user-top-view-office-keyboard-154199/>
- <https://pixabay.com/en/brain-cognition-design-art-2029391/>
- <https://pixabay.com/en/computer-workstation-server-monitor-158743/>
- <https://pixabay.com/en/cpu-processor-intel-amd-chip-152656/>
- <https://pixabay.com/en/microprocessor-processor-cpu-chip-152599/>
- <https://pixabay.com/en/computer-cyber-circuitry-circuits-3163436/>
- <https://pixabay.com/en/ram-computer-memory-module-148579/>