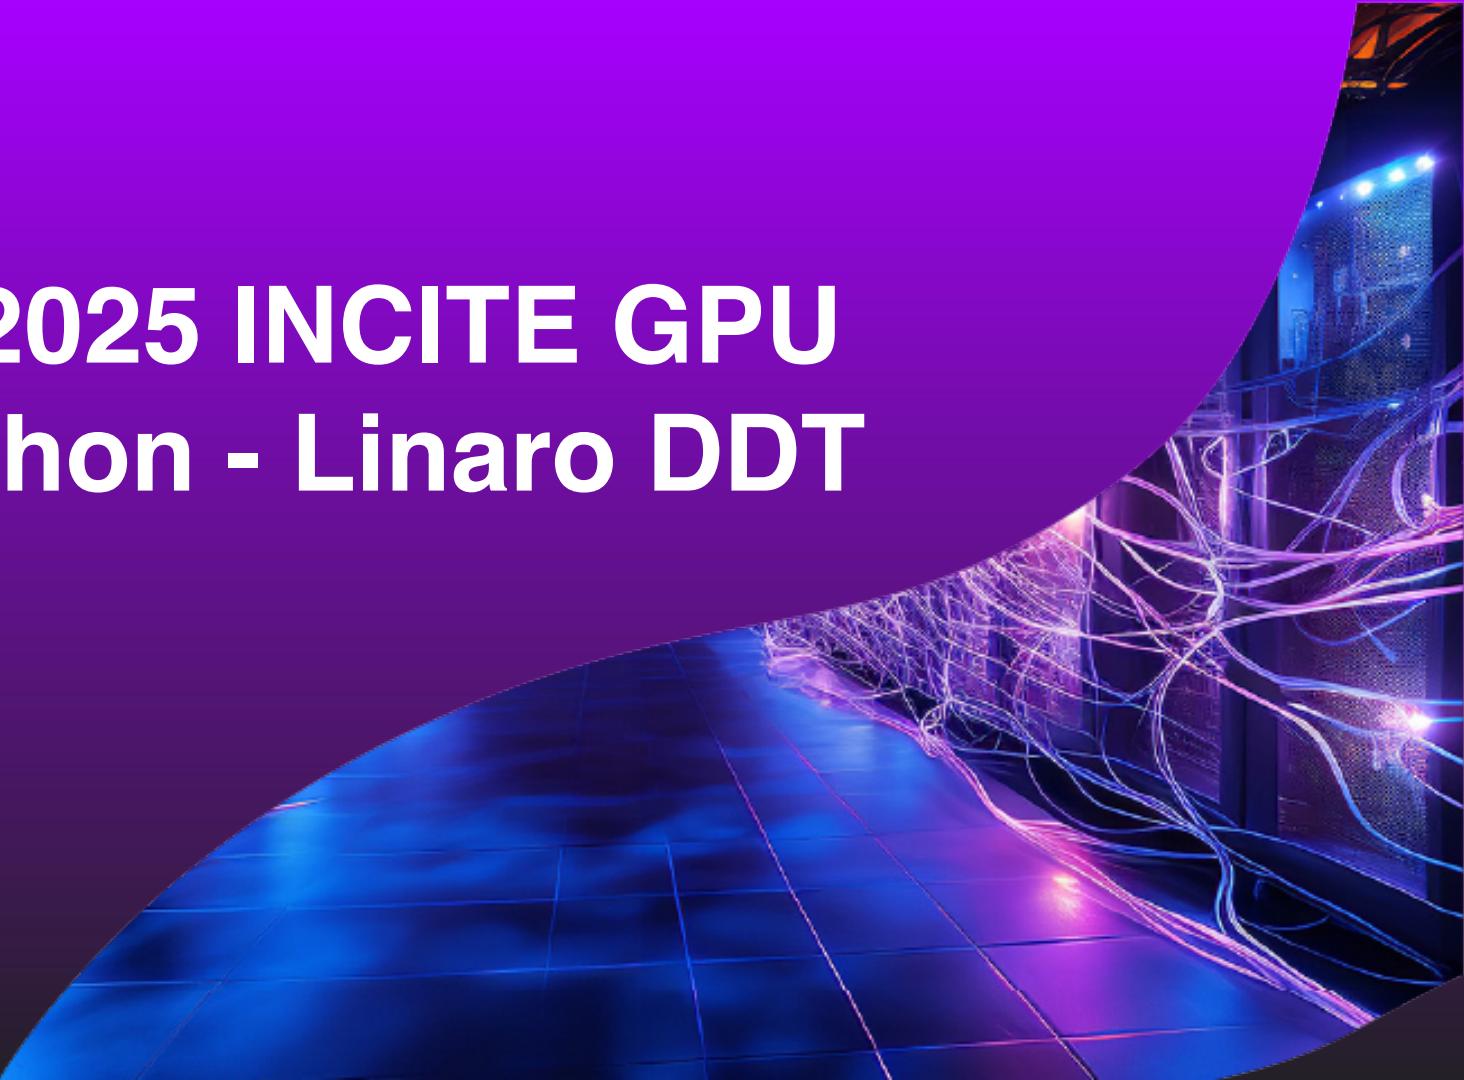


linaroforge

# ALCF 2025 INCITE GPU Hackathon - Linaro DDT

Rudy Shand  
Principal FAE



# Supported Platforms



# Linaro DDT Debugger Highlights

The scalable print alternative

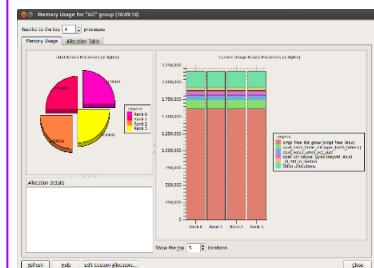
Stop on variable change

Static analysis warnings on code errors

## Memory Debugging

```
# Enable reading of debug environment variables:  
host:~/demo> export NEOReadDebugKeys=1  
  
# Disable RTLD_DEEPBIND, so that malloc/free calls bind to  
# our memory debug library, instead of glibc:  
host:~/demo> export NEODisableDeepBind=1
```

Detect read/write beyond array bounds



# Terminology

	NVIDIA GPU	AMD GPU	Intel GPU
<b>Name</b>	CUDA	ROCM	Xe
<b>Language</b>	CUDA C/C++ and Fortran	HIP	SYCL
<b>Execution Unit</b>	Warp	Wavefront	Sub-group
<b>EU Size</b>	32 Threads	64 Threads	8, 16 or 32 Threads
<b>EU GDB</b>	...	GDB Thread	GDB Thread
<b>EU Thread</b>	...	Lane	Lane (work-item)
<b>Forge GPU Thread</b>	Lane	Lane	Lane

# DDT UI

- 1 Process controls
- 2 Process groups
- 3 Source Code view
- 4 Variables
- 5 Evaluate window
- 6 Parallel Stack
- 7 Project files
- 8 Find a file or function

The screenshot illustrates the DDT UI interface with numbered callouts pointing to various features:

- 1** Process controls: Shows a summary of processes (512 total), current selection (1 process on comp000, pid 1003), and process groups (Group 1 and Group 2).
- 2** Process groups: A list of process groups with their counts and details.
- 3** Source Code view: Displays the C code for `hello.c`, specifically the MPI communication logic.
- 4** Variables: Shows the current variable state in the Locals pane, including `argc`, `argv`, `beingWatched`, etc.
- 5** Evaluate window: Shows the Evaluate pane with expressions like `bigArray[3]`, `my_rank`, and `x + y`.
- 6** Parallel Stack: Shows the Stacks (All) tab with the main thread's stack trace.
- 7** Project files: Shows the Project Files pane with the Application Code tree, currently expanded to show `hello.c`.
- 8** Find a file or function: Shows the search bar and results for `hello.c`.

```

    // Greetings from process %d!, my_rank);
    printf("sending message from (%d)\n", my_rank);
    dest = 0;
    /* Use strlen(message)+1 to include '\0' */
    MPI_Send(message, strlen(message) + 1, MPI_CHAR, dest, tag, MPI_COMM_WORLD);
    beingWatched--;
}
else {
    /* my_rank == 0 */
    for (source = 1; source < p; source++) {
        printf("waiting for message from (%d)\n", source);
        MPI_Recv(message, 100, MPI_CHAR, source, tag, MPI_COMM_WORLD, &status);
        printf("%s\n", message);
        beingWatched++;
    }
    for (i = 1; i < argc; i++)
        if (argv[i] && !strcmp(argv[i], "memcrash"))
            func3();
    for (i = 1; i < argc; i++)
        if (argv[i] && !strcmp(argv[i], "guardafter"))
            func3();
}
    
```

# Debugging Intel Xe GPUs

## Using Linaro DDT

Debug code simultaneously on the GPU and the CPU

Controlling the GPU execution:

- All active threads in a Sub-group will execute in lockstep. Therefore, DDT will step 16 threads at a time.
- Play/Continue runs all GPU threads
- Pause will pause a running kernel

Key (additional) GPU features:

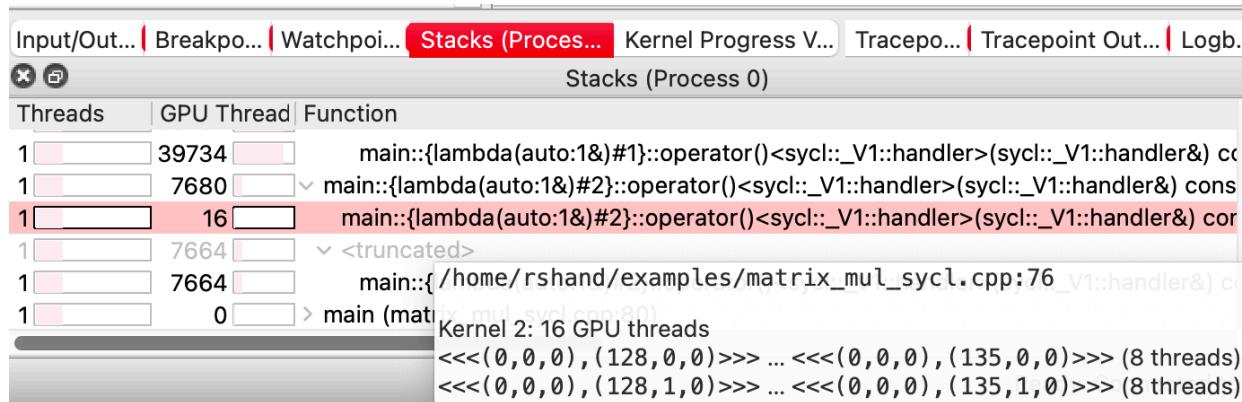
- Kernel Progress View
- GPU thread in parallel stack view
- GPU Thread Selector
- GPU Device Pane

Kernels must be compiled with the `-g` and `-O0` flags

The screenshot shows the Linaro DDT interface with several panes:

- Top Bar:** Shows GPU device selection (Intel(R) Data Center GPU Max 1550), current rank (0-61), and core count (448).
- GPU Threads Tab:** Displays a grid of 64 GPU threads (0-63). A red circle highlights thread 61.
- Project Files Tab:** Shows the file structure of the project, including files like indexer.c, init.c, ion.c, kdreg2\_mem\_monitor.c, libgcc2.c, log.c, and main.c.
- Code Editor Tab:** Displays the C++ code for the `matrix_mul_sycl.cpp` file. The code performs matrix multiplication using OpenCL command groups. A red circle highlights line 61 of the code.
- Kernel Progress View Tab:** Shows the progress of the kernel execution across three work items (M, N, P). The progress bar is nearly complete for M and N, while P is still in progress. A legend indicates that green bars represent scheduled threads and blue bars represent selected threads.
- GPU Devices Tab:** Shows the current GPU configuration with Ranks 0-61, 1 ID, and 448 Cores.

# Parallel Stack View



- Display location and number of threads
- Click item:
  - Select GPU Thread.
  - Update variable display.
  - Move Source Code Viewer.
- Tooltip displays:
  - GPU Thread Ranges.
  - Size of each range.

# Python Debugging

- Debug Features
  - Sparklines for Python variables
  - Tracepoints
  - MDA viewer
  - Mixed language support
- Improved Evaluations:
  - Matrix objects
  - Array objects
  - Pandas DataFrame
  - Series objects
- Python Specific:
  - Stop on uncaught Python exception
  - Show F-string variables
  - Mpi4py, NumPy, SciPy

Linaro DDT - Linaro Forge 23.1

Current Group: All Focus on current: Group Process Thread Step Threads Together

All 0 1 2 3 4 5 6 7

Create Group Project Files

Search (⌘K)

mmult...

```

124     if mr == 0:
125         if fortran_style_array_order:
126             mat_a = numpy.ndarray(shape=(sz, sz), dtype='d', order='F')
127             mat_b = numpy.ndarray(shape=(sz, sz), dtype='d', order='F')
128             mat_c = numpy.ndarray(shape=(sz, sz), dtype='d', order='F')
129         else:
130             mat_a = numpy.ndarray(shape=(sz*sz), dtype='d', order='C')
131             mat_b = numpy.ndarray(shape=(sz*sz), dtype='d', order='C')
132             mat_c = numpy.ndarray(shape=(sz*sz), dtype='d', order='C')
133
134     print("(): Initializing matrices...".format(mr))
135     minit(sz, fortran_style_array_order, mat_a)
136     minit(sz, fortran_style_array_order, mat_b)
137     minit(sz, fortran_style_array_order, mat_c)
138
139     print("(): Sending matrices".format(mr))
140     for i in range(1, nproc):
141         # Get a slice from the mat_a and mat_c matrix
142         if fortran_style_array_order:
143             mat_a_slice = mat_c[:, i*mslice_r:(i+1)*mslice_r]
144             mat_c_slice = mat_c[:, i*mslice_r:(i+1)*mslice_r]
145

```

Input/Output | Breakpoints | Watchpoints | Stacks (All) | Tracepoints | Tracepoint Output | Logbook | Evaluate

Processes | Function

Stacks (All)

Name | Value

mslice	512
nproc	8

1 <module> (allinea\_ddt\_trace.py:155)

1 main (allinea\_ddt\_trace.py:140)

1 <module> (mmult.py:215)

1 main (mmult.py:134)

```
ddt --connect mpiexec -n 8 python3
%allinea_python_debug% ./mmult.py
```

# Run DDT in offline mode

Run the application under DDT and halt or report when a failure occurs

You can run the debugger in non-interactive mode

- For long-running jobs / debugging at very high scale
- For automated testing, continuous integration...

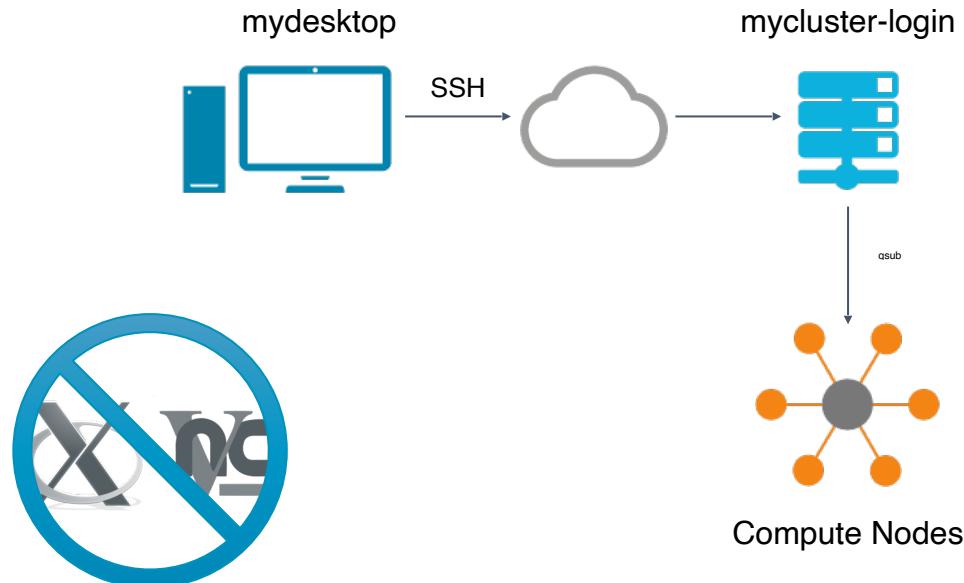
To do so, use following arguments:

- \$ ddt --offline --output=report.html mpirun ./jacobi\_omp\_mpi\_gnu.exe
  - --offline enable non-interactive debugging
  - --output specifies the name and output of the non-interactive debugging session
    - Html
    - Txt
  - Add --mem-debug to enable memory debugging and memory leak detection

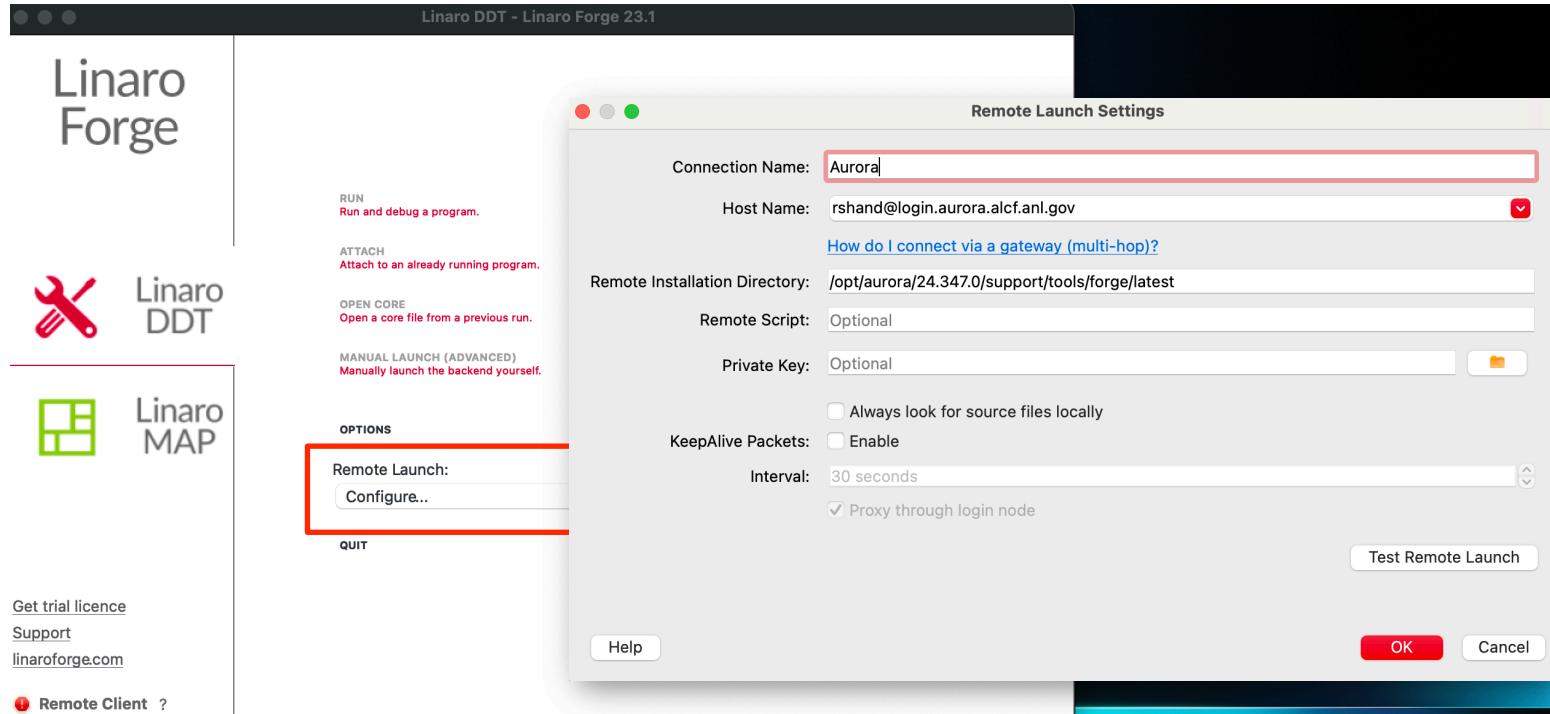
```
ddt --offline -o jacobi_omp_mpi_gnu_debug.txt \
        --trace-at _jacobi.F90:83,residual \
        mpiexec ./jacobi omp mpi gnu.exe
```

# The Forge GUI and where to run it

DDT provides a powerful GUI that can be run in a variety of configurations.

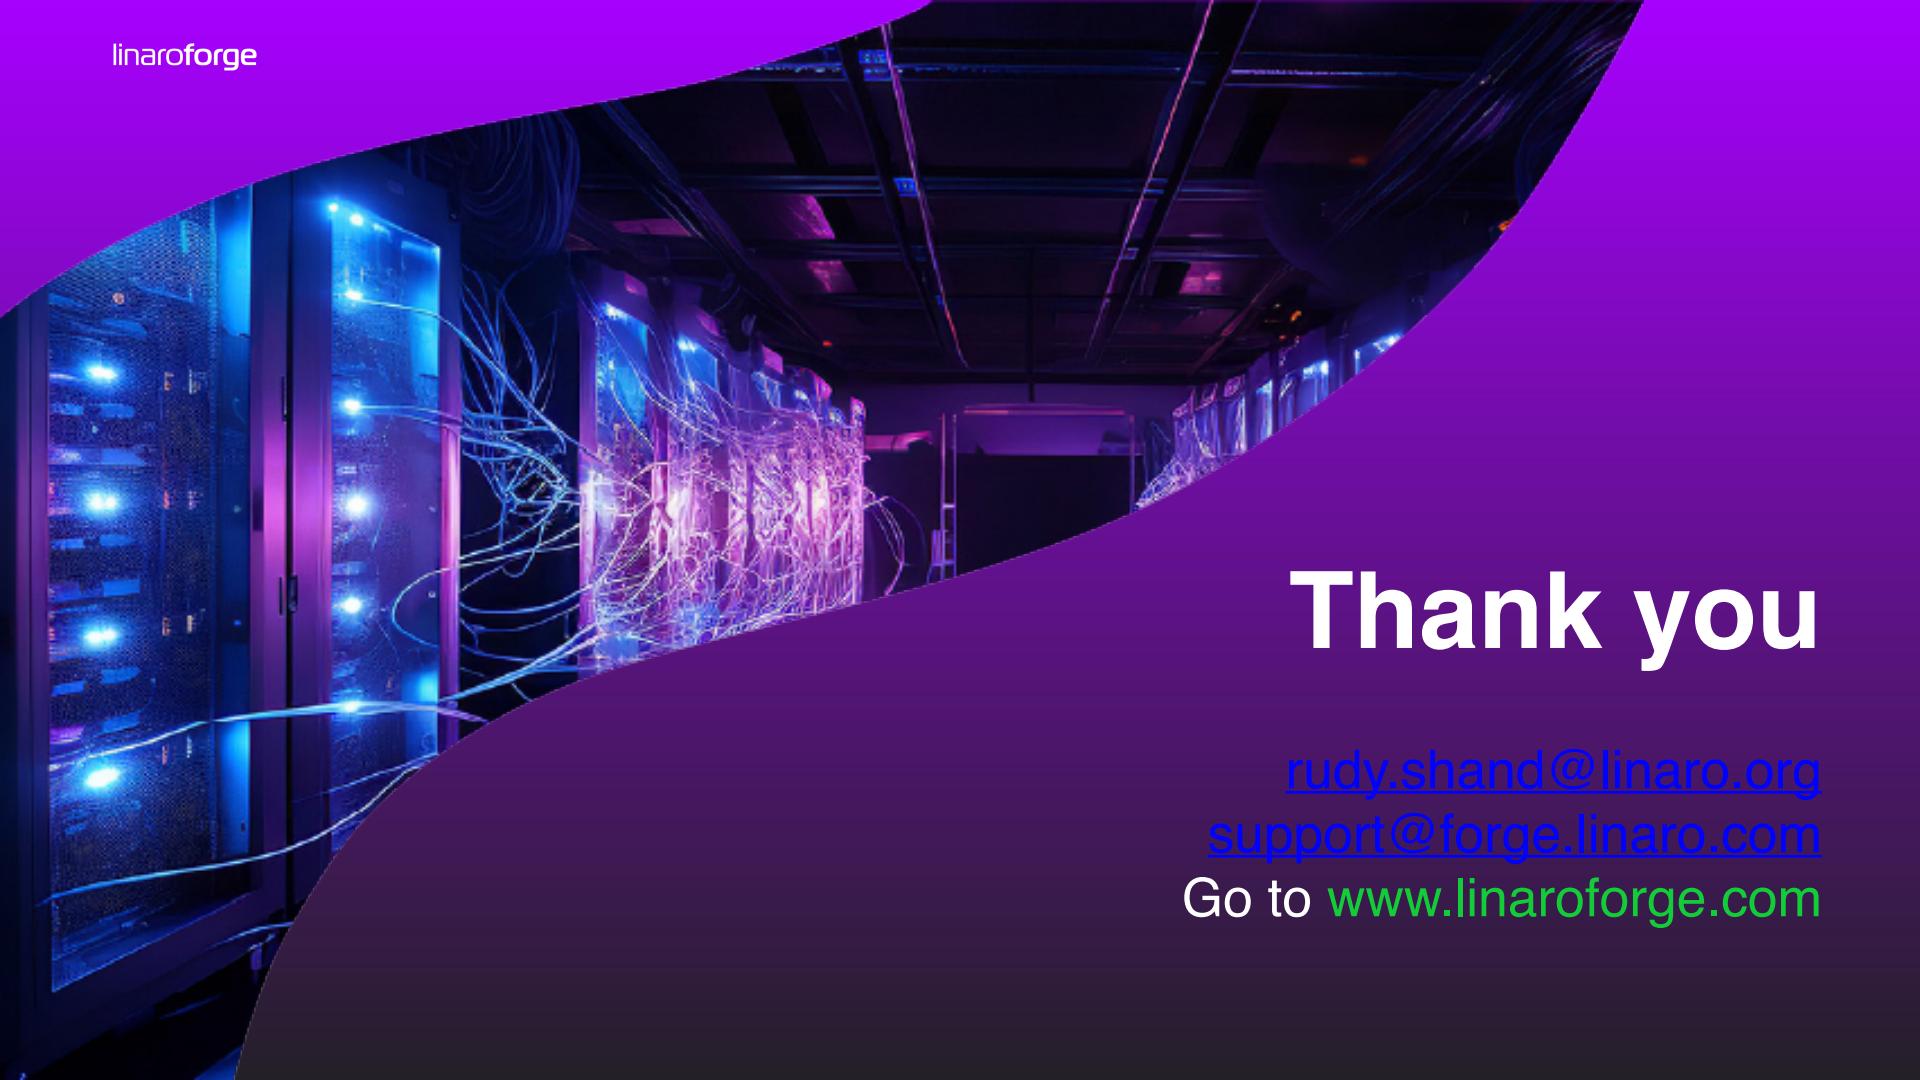


# Remote connection to AWS



# Debug with DDT on Aurora

```
mpicxx -fsycl -g -O0 matrix_mul_sycl.cpp -o matrixmul  
qsub -l select=2 -l walltime=30:00 -l filesystems=flare -A gpu_hack -q gpu_hack_prio -I  
. /soft/compilers/oneapi/2025.1.0/debugger/2025.1/env/vars.sh  
https://docs.alcf.anl.gov/aurora/debugging/ddt-aurora/#invoking-the-ddt-server-from-aurora  
ddt --np=24 --connect --mpi=generic --mpiargs="--ppn 12 --envall" ./matrixmul
```



# Thank you

[rudy.shand@linaro.org](mailto:rudy.shand@linaro.org)

[support@forge.linaro.com](mailto:support@forge.linaro.com)

Go to [www.linaroforge.com](http://www.linaroforge.com)