

Home

What is FAbRIC

Derek Chiou (PI, UT Austin/Microsoft)

System admins: **Alex Hsu** (UT Austin), **David Carver** (UT Austin) and **Laura Timm** (UT Austin)

Reconfigurable hardware has tremendous potential in terms of performance and power efficiency. However, the high cost of such systems, the high cost of the tools needed to program such systems, and need to program at a very low level together put such systems out of reach of most researchers. The FAbRIC (FPGA Research Infrastructure Cloud) project will acquire and maintain such systems and their tools in the [Texas Advanced Computing Center](#) for open use by all researchers. In addition, the PIs will port our own reconfigurable hardware "operating systems" to the platforms to dramatically improve their usability and our own applications to serve as starting points for future work.

Having an open, shared resource of this kind makes some of the highest performance computational power available to all researchers, regardless of their location or their ability to purchase and maintain such systems. The project team will run distributed classes for students and researchers to learn how to use such platforms. In addition, the default usage model of "open source to play" enables anyone who agrees to let the team publish their open source code to use the facility free of charge. Such a shared platform with free industrial-strength tools and open sourced code finally enables true reproducibility of research results and the ability to leverage others' work.

The award started June 1, 2012. So far we have deployed two systems, **Convey MX** and **IBM POWER8+CAPI**. The third system, **Microsoft Catapult**, is being brought up and will be available to beta testers soon. All these platforms are equipped with FPGA accelerators and development servers in a cloud based production environment. To be available for open use, FAbRIC systems are placed in the Texas Advanced Compute Centre (TACC), UT Austin's supercomputer center.

In addition to NSF support, Intel (formerly Altera), Nallatech, [Xilinx](#) and Alpha-data have all committed FPGA/board donations. Intel has committed their entire suite of CAD tools, IBM has donated POWER8 servers, Intel has provided funds for the CAD tool servers, Microsoft has provided Catapult servers and funds for operating them, Nvidia has donated their GPUs, [Bluespec](#) has committed their Bluespec compiler, and [Impulse Accelerated Technologies](#) has committed their ImpulseC C-to-gates compiler. We are in active discussions with other companies who are interested in contributing their technologies to FAbRIC.

This material is based on work supported by the National Science Foundation under Grant No. 1205721 and generous donations and technical support from Alpha-data, Intel, Bluespec, IBM, ImpulseC, Intel, Microsoft, Nallatech, Nvidia and Xilinx.

In any publications that use FAbRIC, please include the following acknowledgment: "This material is based on work supported by the National Science Foundation under Grant No. 1205721."

User Access

If you would like to use the FAbRIC system, **please first request and obtain a TACC account at <https://portal.tacc.utexas.edu/account-request>**. Make sure you get a confirmation email from TACC and that you can login to the user portal. Keep in mind you will need to login to your TACC account at least once to activate your account.

Forward that confirmation email to account-request@openfabric.org and the system admin of the platform you'd like to use, along with a brief reason why you want the account. Please provide your full legal name, your citizenship, and your current residential address, and your university name. You should send the email from your university email address.

In addition, include the following statement in your email.

- "All of the code that I will pass through FAbRIC CAD tools (such as Verilog files, Bluespec files, etc.) and the files needed to process that code (such as Makefiles) is either already open source (GPL version 2 or above, BSD, or MIT licenses) or I have the right to make it open source and are hereby making all of the code that I pass through FAbRIC CAD tools open source by one of those licenses. I will provide access to my source code to the CAD tool vendors and the FAbRIC administrators immediately. The simplest way to do that is to provide a repository account to the FAbRIC administrators. By default, the CAD tool vendors and/or the FAbRIC administrators agree not to publish the code publicly for at least 12 months.

I acknowledge that the tools, servers, and FPGAs are potentially subject to export controls under U.S. and other applicable Government laws and regulations. I will comply with these laws and regulations and agree to obtain all required Government authorizations.

In any publications that use FAbRIC, I agree to include the following acknowledgment: [[[This material is based on work supported by the National Science Foundation under Grant No. 1205721.]]]"

If you want access to Microsoft Catapult, add the following to your email.

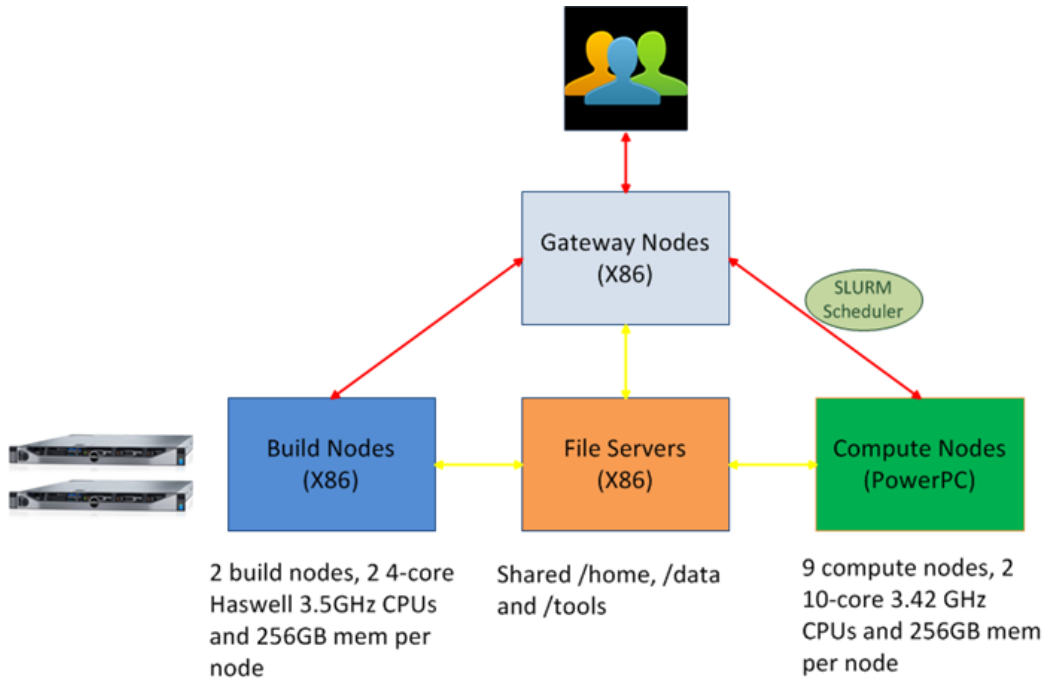
- "I acknowledge that my access to and use of Microsoft 's Project Catapult Academic Shell and Driver and related hardware provided by Microsoft is governed by, and subject to, the terms and conditions of the Microsoft Research License Agreement for the [Microsoft Project Catapult Academic Shell and Driver](#). By accessing or using Microsoft Project Catapult materials, I represent and warrant that I have read the agreement, and I agree to be bound by it."

If you would like access to Intel FPGAs and tools (IBM Power8+CAPI, Microsoft Catapult), please also get a myAltera account (<https://www.altera.com/mal-all/mal-signin.html>) and forward the confirmation email to account-request@openfabric.org as well.

You will receive an email once your FAbRIC account is approved and becomes active.

IBM POWER8+CAPI cluster

The IBM POWER8+CAPI Cluster is a cluster of several x86 servers and nine POWER8 servers. Each POWER8 node is a heterogeneous platform capable of running GPGPU- and/or FPGA- accelerated applications.



Our current setup supports three accelerating devices:

- Nallatech 385 A7 Stratix V Altera FPGA adapter
- Alpha-data 7V3 Virtex7 Xilinx-based FPGA adapter
- NVIDIA Tesla K40m GPGPU card

Currently each POWER node has one node from each vendor. FPGA boards are IBM CAPI (Coherent Accelerator Processor Interface) enabled to provide coherent shared memory between the processor and accelerators.

The system is now available to the public. Our early users have been able to generate promising results from running real-world applications.

A quick user start guide: <https://docs.google.com/document/d/1EkWfSXGTOqAwsil2pzNj99ovmzB9d-ad-C6YrOaFvDo/edit#heading=h.qf46aqht3ic>

Primary system admins: Alex Hsu (alex.hsu@tacc.utexas.edu) and David Carver (dcarver@tacc.utexas.edu)

Microsoft Catapult cluster

The Microsoft Catapult system consists of 432 two-socket Intel Xeon-based nodes, each with 64 GB of memory and an Altera Stratix V D5 FPGA with 8 GB of local DDR3 memory. FPGAs communicate to their host CPUs via a PCIe Gen3 x8 connection, providing 8GB/s guaranteed-not-to-exceed bandwidth, and each FPGA can read and write data stored on its host node using this connection. The FPGAs are connected to one another via a dedicated network using high-speed serial links. This network forms a two dimensional 6x8 torus within a pod of 48 servers, and provides low latency communication between neighboring FPGAs. This design supports the use of multiple FPGAs to solve a single problem, while adding resilience to server and FPGA failures.

Per Node:

- Two Intel® Xeon® E5-2450, 2.1GHz, 8-core, 20MB Cache, 95W
- 64GB RAM
- Four 2TB 7.2k 3G SATA 3.5"; Two 480GB 6G Micron SATA SSD 2.5"
- Intel 82599 10GbE Mezz Card

- Altera Stratix V FPGA Card
- Operating System: Windows Server 2012



Convey MX system

The Convey MX system consists of a gateway node (login1.fabric.tacc.utexas.edu) and a compute node (c0-1.fabric.tacc.utexas.edu). You will need to login to the gateway node in order to ssh into the compute node.

The gateway node: Dell R720 server, 64GB memory, 16 cores (Intel® Xeon® CPU E5-2670 @ 2.60GHz)

The compute node is a Convey MX system, with 128GB of RAM, 100GB of 64b granularity bandwidth, and 4 user "application" FPGAs.

There are two filesystems, /home and /data that are mounted on both login1 and c0-1. /home is backed up however /data is not. \$HOME and \$DATA are set for your convenience. The command cdd will cd to your /data directory.

Convey packages are installed on both nodes. Convey's documentation, including a Programmer's Guide, is at http://www.conveysupport.com/help/?page_id=112 (free registration on Convey support website required).

Note: The Convey MX system is being phased out due to old hardware and low usage.

Primary system admins: Alex Hsu (alex.hsu@tacc.utexas.edu) and David Carver (dcarver@tacc.utexas.edu)

History News

News (November 10, 2015): The second FABRIC platform is IBM POWER8+CAPI. We have nine of those systems that are being brought up right now. We will be opening up the systems to a limited number of beta testers in the next month or two, and plan to open up to the general research community in the first quarter of next year. Thanks to IBM for donating the POWER8 servers, Altera/Nallatech for donating their FPGA boards, Nvidia for donating their GPU boards, and Xilinx/Alpha Data for donating their FPGA boards.

News (November 12, 2015): The third system type is the Microsoft Catapult platform. 384 data center servers, each equipped with an Altera Stratix V D5 FPGA are in TACC being brought up right now. We will be opening up the systems to alpha users in the next month or two, and plan to open up to beta testers early next year. Thanks to Microsoft and Altera for providing the systems, FPGAs, and tools. Here is a blog post. http://blogs.msdn.com/b/msr_er/archive/2015/11/12/project-catapult-servers-available-to-academic-researchers.aspx