

# The ARM ecosystem for scientific computing

Geraint North  
Architect, High Performance Ecosystem  
ARM Ltd



The Architecture for the Digital World®

**ARM**

# In this presentation...



**ARM Technology**



**HPC Partnerships**



**Software Availability**

The ARM ecosystem for scientific computing

# ARM Technology



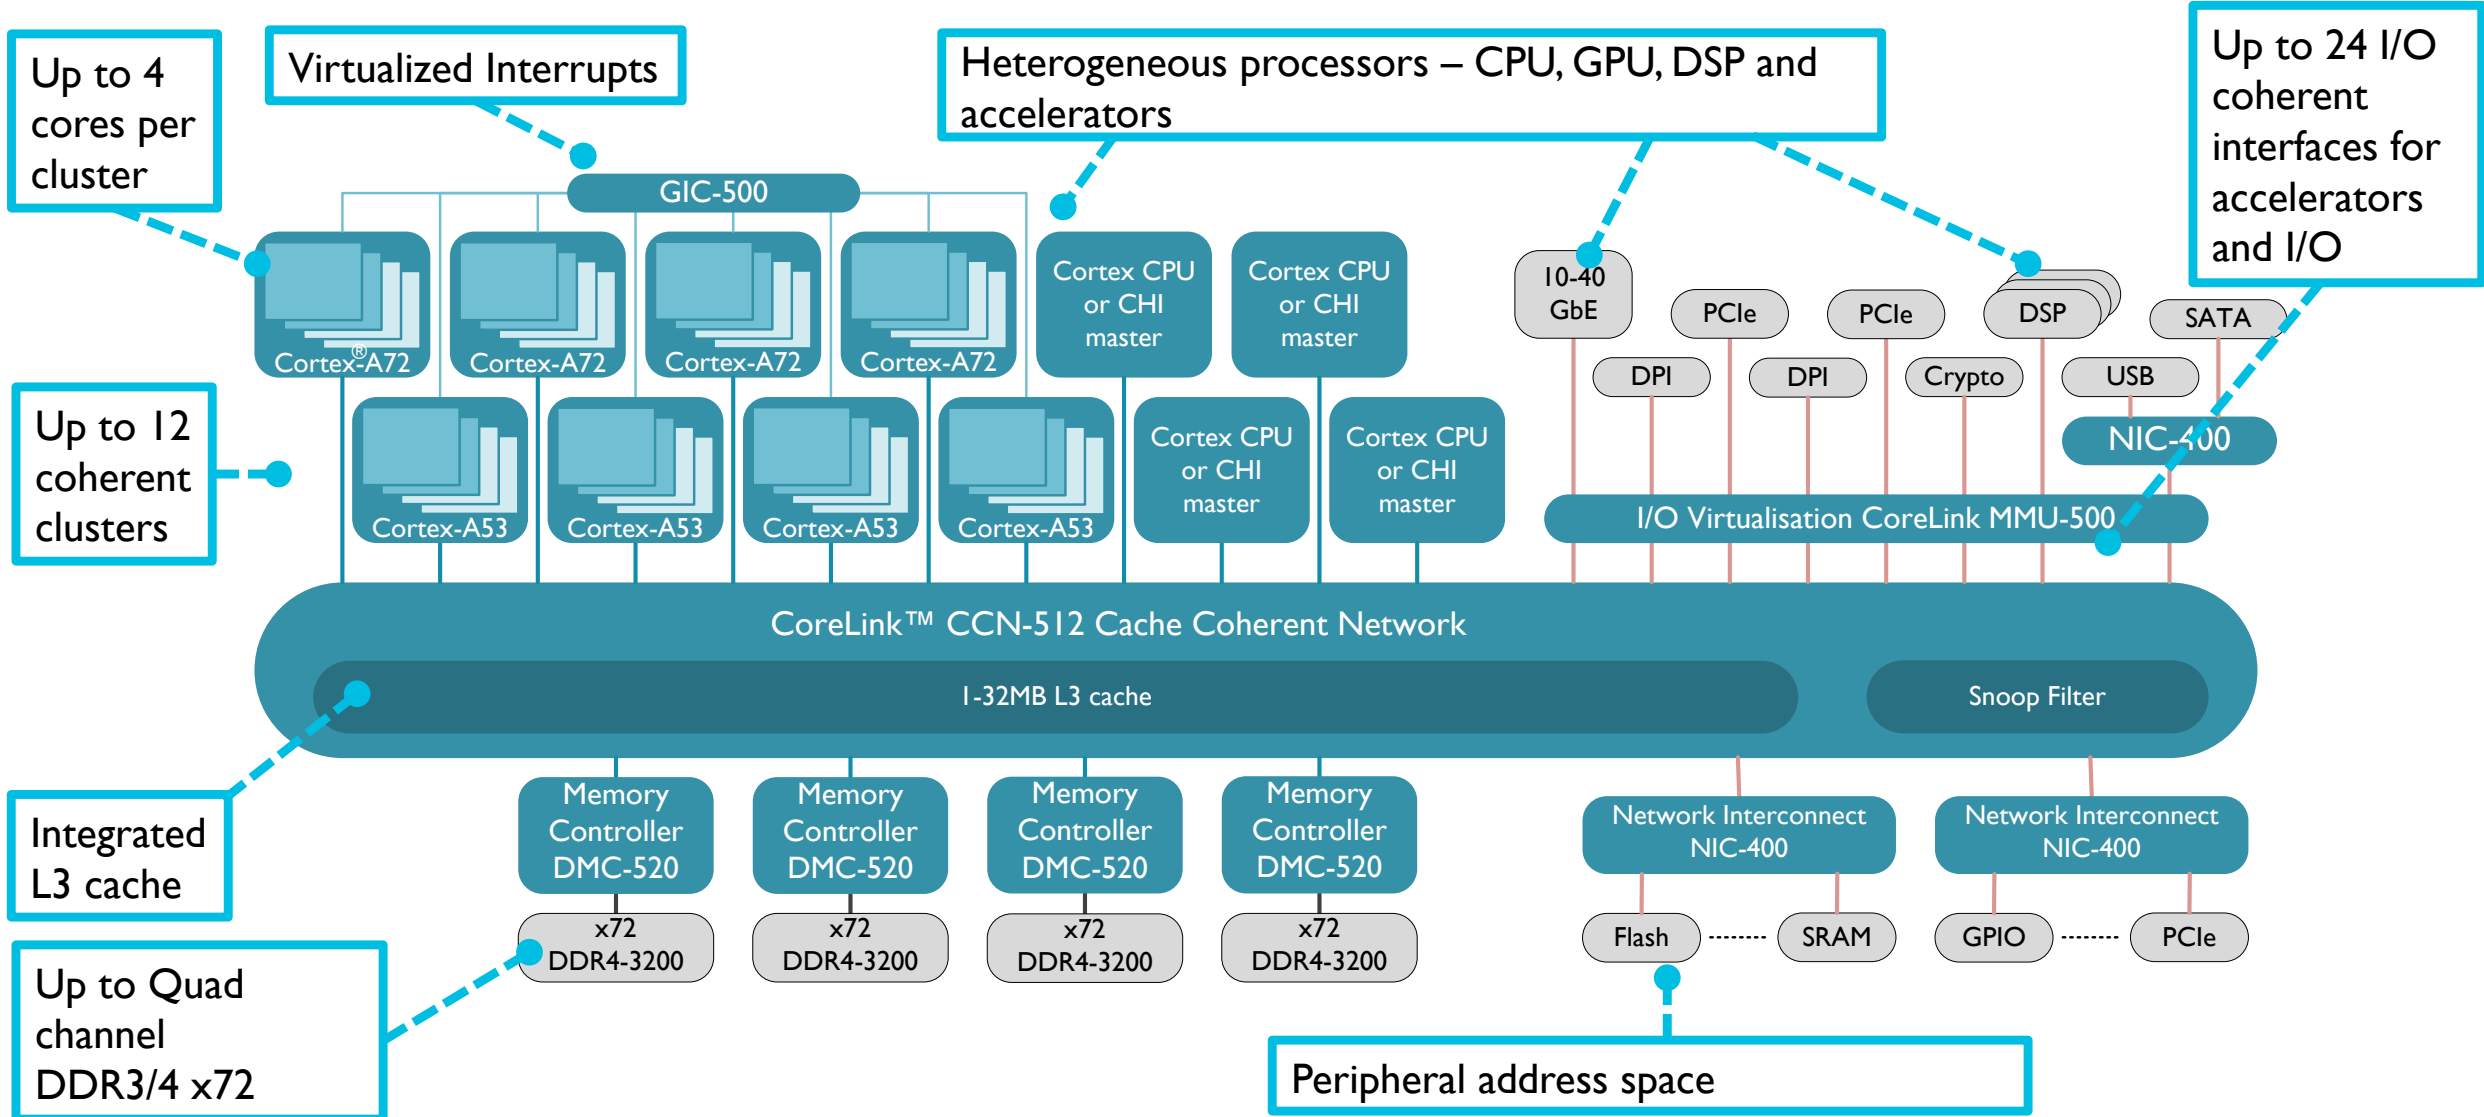
# A brief introduction to ARM

**ARM designs scalable, energy efficient-processors and related technologies to deliver the intelligence in applications ranging from sensors to servers, including smartphones, tablets, enterprise infrastructure and the Internet of Things.**

## **Our business model**


- Our technology is licensed to our partners (currently over 350), who use it to create their own system-on-chip (SoC) products.
- 60 billion chips shipped since ARM was founded in 1990.
- Our technology is present in 95% of smart phones, 80% of digital cameras, and 35% of all electronic devices.
- We may license an instruction set **architecture** (e.g. “ARMv8”) or a specific microarchitectural **implementation** (e.g. “A72”).

# Extensible Architecture for Heterogeneous Multi-core Solutions



# The introduction of 64-bit

- ARM's first 64-bit architecture (ARMv8-A) was announced in 2011.
- Adoption of 64-bit in the mobile segment has accelerated the availability of implementations and software ecosystem support.
- Many popular Linux distributions support 64-bit ARMv8:
  - Debian, Fedora, Red Hat, CentOS, OpenSUSE.
- Many partners now developing 64-bit server-class silicon:
  - Applied Micro X-Gene 1 and 2
  - AMD Seattle
  - Broadcom Vulcan
  - Cavium ThunderX
  - HiSilicon D02
  - watch this space...






# ARM in Datacenter/HPC Compute

- ARM-based SOC's enable compelling solutions:
  - Optimized performance/watt, performance/RU
  - Workload optimized balance of CPU, memory, cache, and IO
  - Application specific HW accelerators
  - Heterogeneous compute (DSP's, FPGA's, GPU's, etc).
  - Comprehensive SW Ecosystem
- Built to existing datacenter standards
  - ARM platforms use standard motherboard form-factors or chassis/rack form-factors
  - Standard interconnects (PCIe, RapidIO, etc.)
  - Standard platform firmware abstractions for deployment and management
  - ARM technology is available in enterprise form factors from **HP (Moonshot), E4 Computing** and **SoftIron**



softIron

APM6408 SERIES



ARM

## FEATURES

**Form Factor:** 2U

**CPU** 1x CAVIUM Thunder-X 48 cores

**GPU up to** 1x NVIDIA Kepler® K20, K40, K80

**Memory** Up to 128 GB RAM DDR3

**Network** 2x 10 GbE SFP+, 1x IB FDR QSFP, 1x 40 GbE QSFP

**Storage** 4x SATA 3.0

**Expansion slots** 1x PCI-E 3.0 x8 (in x16), 1x PCI-E 3.0 x8



## FEATURES

**Form Factor:** 2U

**CPU** 1x APM X-Gene 8 cores

**GPU up to** 2x NVIDIA Kepler® K20, K40, K80

**Memory** Up to 128 GB RAM DDR3

**Network** 2x 10 GbE SFP+, 1x Infiniband FDR QSFP

**Storage** 4x SATA 3.0

**Expansion slots** 2x PCI-E 3.0 x8 (in x16)





# ARM in Enterprise Networking

- 10% share of networking infrastructure equipment, and growing rapidly.
- Deployed in 4G basestations and WLAN infrastructure
  - All based on 32-bit ARMv7 chips
- Major semiconductor networking providers have announced ARMv8-A based chips for this market.
- HiSilicon started manufacturing 32x Cortex-A57 chip on TSMC 16FF process
- *OpenDataPlane*: An open-source, cross-platform set of application programming interfaces (APIs) for the networking data plane.



# ARM Manchester Design Centre

- Founded July 2014
- 15 engineers focused on runtime, library and compiler development for high-performance computing and server.
- Focus areas:
  - Autovectorization improvements for open-source compilers.
  - ARM-tuned numerical libraries.
  - Future ARM architectures
  - OpenMP supported and tuned for ARM.
  - Flang – Fortran for LLVM.



The ARM ecosystem for scientific computing

# HPC Partnerships



# Mont-Blanc

## Objectives

To develop a full energy-efficient HPC prototype using low-power commercially available embedded technology.

To develop a portfolio of exascale applications to be run on this new generation of HPC systems.

To design a next-generation HPC system together with a range of embedded technologies in order to overcome the limitations identified in the prototype system.


## Status

Prototype operational:  
8 standard BullX chassis, 72 compute blades, 1080 compute cards, 2160 ARM Cortex-A15 processors, 1080 ARM Mali-T604 GPUs.



|| Scientific applications ported and in use for evaluation of the prototype.

Research ongoing into areas such as memory, on-chip and off-chip interconnect, compute acceleration



MB Prototype installed in the Torre Girona chapel @ BSC



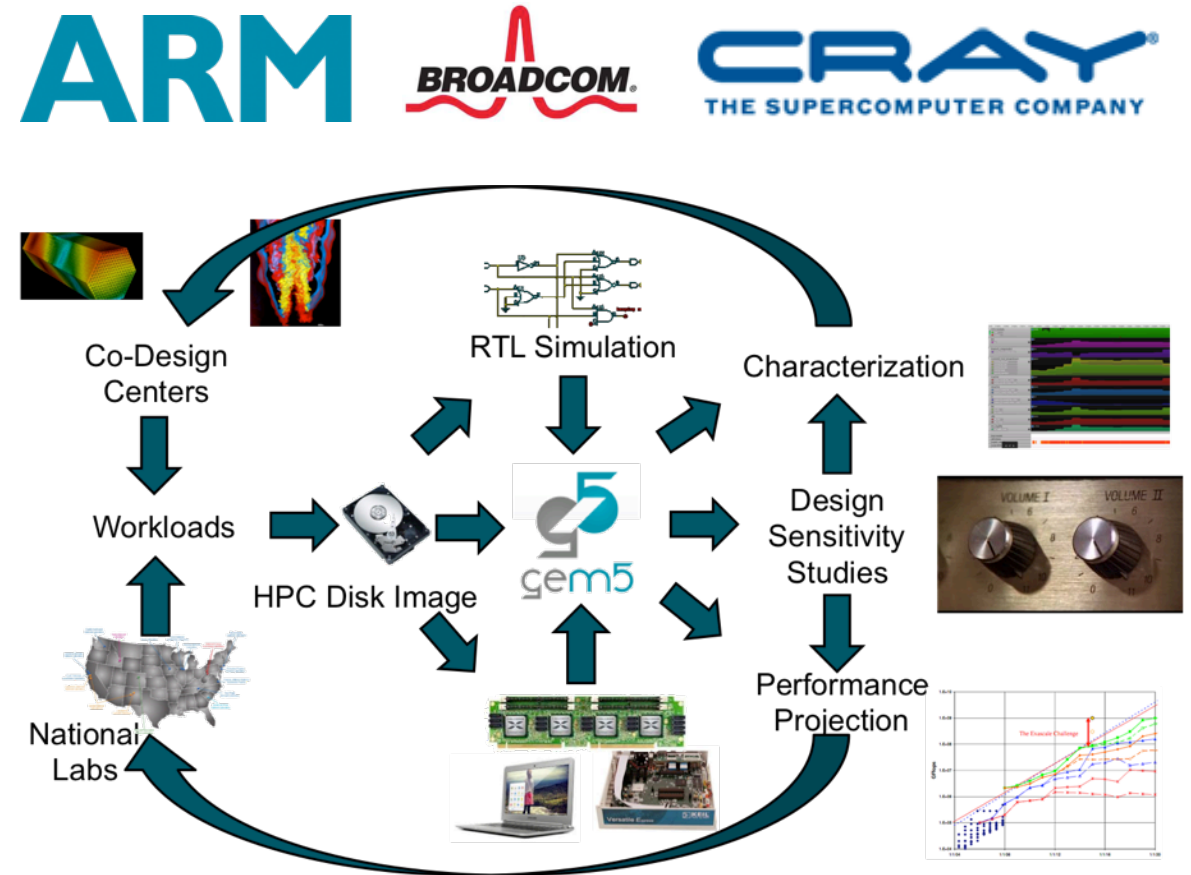
# U.S. Department of Energy: FastForward II

<https://asc.llnl.gov/fastforward/>



## ARM Focus Areas

- Evaluation of next-generation architecture in the context of DoE applications
- Evaluation of throughput and multithreaded core designs for HPC
- Next generation memory technologies
- Design study to find right balance of core types, memory, and interconnect
- Development and integration of full system simulation technology with other partners
- Workload characterization and optimization for ARM architecture





# The Hartree Centre

Science & Technology Facilities Council, UK

The **Energy Efficient Computing Research Programme** has been established through a £19M capital grant from the Department of Business Innovation and Skills to establish a centre of best practice in the UK that will enable users of computer systems to achieve the same outcomes while minimising the consumption of energy.



Lenovo are providing a NeXtScale system: 1,152 64-bit Cavium ThunderX ARM cores in 6U.




Science & Technology  
Facilities Council

*“This is a fantastic opportunity to meet the challenge of developing a computationally powerful and energy-efficient platform based on the 64-bit ARM v8 microprocessor ... The Hartree centre will be actively developing a robust software ecosystem encompassing compilers, linkers, numerical libraries and tools – all of which are fundamental to the adoption of these types of technologies.”*

# Polly Labs

- **Goal:** Promote the adoption of polyhedral compilation techniques in the LLVM compiler
  - Polyhedral techniques allow nested loops to be restructured to improve parallelism, vectorization potential, data access locality etc.
- Delivers production-quality technology into the LLVM open-source compiler.
- Raises the bar for aggressive loop optimization in open-source compilers, bringing them closer to the capabilities of commercial compilers.
- Open to new partners, looking to mutualize efforts and deliver maximum value to the LLVM project.

Founding Partners



The ARM ecosystem for scientific computing

# Software Availability



# Compilers

## Commercial

### PathScale (Alpha)

C, C++, Fortran  
OpenMP 4.0

November 2014:  
PathScale provides the full EKOPath compiler suite including OpenACC and OpenMP 4.0 C/C++/Fortran support for ARMv8 to support HPC and Enterprise customers exploring the power efficiencies of these devices. As an enabling technology, EKOPath gives our customers the ability to compile for native ARMv8 CPU or accelerated architectures that return the fastest time to solution. Your application defines the benchmark, EKOPath lets you evaluate the new architecture with your code, across either Intel64/AMD64 and now directly compare it against the performance of enterprise ready ARMv8 processors.

### NAG (Alpha)

Fortran  
OpenMP 3.1

November 2013:  
The [Numerical Algorithms Group \(NAG\)](#), the global numerical software and HPC services company, announces a new technical collaboration with ARM®, the world's leading semiconductor IP supplier. NAG's highly skilled team of HPC experts, numerical analysts and computer scientists will ensure the algorithms in the NAG Numerical Library and the facilities of the NAG FORTRAN Compiler are available for use on ARM's 64-bit ARMv8-A architecture-based platforms.

## Open-Source

### **GCC**

C, C++, Fortran,  
OpenMP 4.0

- Open-source focus on AArch64 correctness up to 2014.
- Now improving core performance through mostly architectural (not microarchitectural) optimisations.
- Command-line enablement for new ARM cores (e.g. A72).
- Most focus and improvement in floating-point code.

#### **Current work:**

- Improvements for big-endian ARM.
- Floating-point rounding mode optimization.
- Making use of more sophisticated ARM instructions.
- Scheduler / register allocation improvements.
- Improved memcpy, memset, glibc string routines.
- Improved performance on NEON intrinsics.

### **LLVM**

C, C++  
OpenMP 3.1

#### **Current work:**

- Vectorizer improvement.
- Loop unrolling/interleaving.
- Improved register allocation.
- ABI conformance.
- Improve inliner heuristics.
- Scheduling for Cortex-A57.
- Software pipelining.
- Jump threading.

# Allinea Forge Development Toolsuite

- Development tools for high performance software
  - For multi-threaded and multi-process software
  - C, C++ and F90
- The market leader in technical and parallel computing
- Allinea DDT
  - debug complex software problems
  - Includes memory debugging, visualization, C++ STL, F90 and more
  - Available NOW for ARM 64-bit
- Allinea MAP
  - profiler for optimizing code performance
  - Tackle I/O, threads, CPU, synchronization and memory performance issues
  - Coming soon for ARM 64-bit

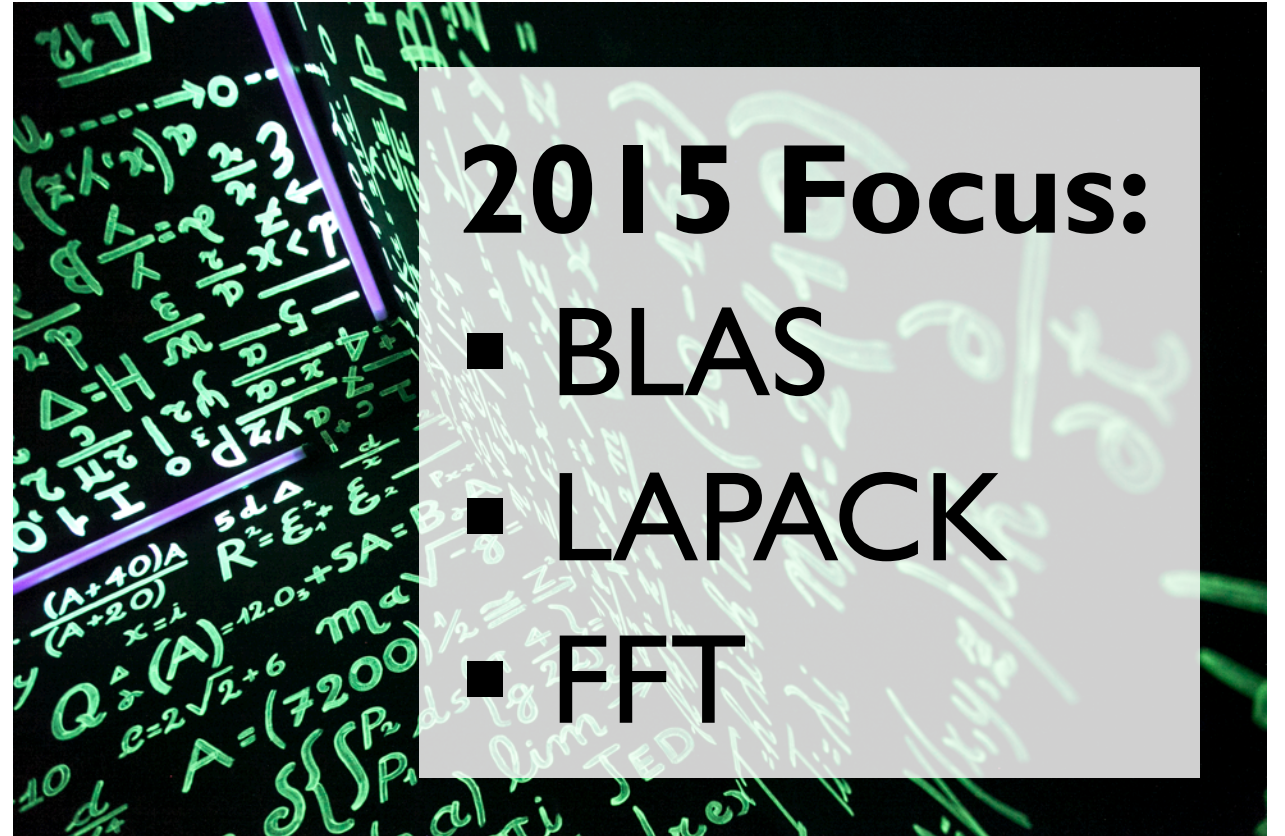




# ARM Math Libraries

In 2015, we plan to offer a commercially supported set of numerical libraries for scientific computing.

By the end of 2015, an HPC-specific ARM microsite will offer downloads, technical reference material, how-to-guides and third-party software recommendations for the scientific computing community.



**2015 Focus:**

- BLAS
- LAPACK
- FFT

# In summary...



## **ARM Technology**

64-bit high-performance ARM cores are available today.



## **HPC Partnerships**

Scientific establishments are deploying and evaluating ARM with real HPC workloads.



## **Software Availability**

Key ecosystem pieces are available and mature.

Thank you

