


Big Data and Extreme Scale Computing (BDEC)

Happy Sithole
CHPC –South Africa
16th July 2015




SKA Core in the Karoo




SCIENCE DATA PROCESSOR



SOUTH AFRICA



WESTERN AUSTRALIA



SKA Work Packages



Work Packages

Assembly, Integration and Verification (AIV)

Central Signal Processor (CSP)

Dish (DSH)

Infrastructure Australia and Africa (INFRA AU/
INFRA SA)

Low-Frequency Aperture Array (LFAA)

Mid-Frequency Aperture Array (MFAA)

Signal and Data Transport (SaDT)

Science Data Processor (SDP)

Telescope Manager (TM)

Wideband Single Pixel Feeds (WBSPF)

Lead/SE

SE

SE, Feeds, Optics, Prototype

Lead/SE

SE, LMC, LINFA, PIP.IMG

SDP consortium members



Management Groupings	Workshare (%)
University of Cambridge (Astrophysics & HPFCS)	9.15
Netherlands Institute for Radio Astronomy	9.25
International Centre for Radio Astronomy Research	8.35
SKA South Africa / CHPC	8.15
STFC Laboratories	4.05
Non-Imaging Processing Team	6.95
University of Manchester	
Max-Planck-Institut für Radioastronomie	
University of Oxford (Physics)	
University of Oxford (OeRC)	4.85
Chinese Universities Collaboration	5.85
New Zealand Universities Collaboration	3.55
Canadian Collaboration	13.65
Forschungszentrum Jülich	2.95
Centre for High Performance Computing South Africa	3.95
iVEC Australia (Pawsey)	1.85
Centro Nacional de Supercomputación	2.25
Fundación Centro de Supercomputación de Castilla y León	1.85
Instituto de Telecomunicações	3.95
University of Southampton	2.35
University College London	2.35
University of Melbourne	1.85
French Universities Collaboration	1.85
Universidad de Chile	1.85

Resourcing




Work Breakdown Structure for the SDP

Units: FTEs (Full-Time Equivalent)																											
L4	Description	ASTRON	BSC	CADC	CHPC	CSIRO	ICRAR	IVEC	JULICH	MPIFR	NZA	PRC	PTC	SKA SA	SOTON	STFC	UCAM	UCL	UCT	UFOP	UFUO	UMAN	UOXF	TOTAL			
MGT	PROJECT MANAGEMENT	0.16	0.00	0.00	0.00	0.00	1.00	0.10	0.00	0.05	0.00	0.00	0.00	0.20	0.00	0.00	2.00	0.00	0.00	0.00	0.00	0.00	0.15	0.00	3.66		
SE	SYSTEM ENGINEERING	0.00	0.00	0.00	0.00	0.00	2.00	0.00	0.00	0.00	0.00	0.25	0.00	3.00	0.00	0.00	0.50	0.00	0.00	0.00	0.00	0.00	0.00	0.00	5.75		
ARCH	Architecture Definition, Design & Analysis	0.00	0.00	0.00	0.00	0.00	1.00	0.00	0.50	0.00	0.00	0.45	0.00	0.00	0.00	2.10	2.75	0.00	0.00	0.00	0.00	0.00	0.05	6.85			
COMP	COMPUTING PLATFORM	0.92	0.50	0.00	0.00	0.00	0.75	0.40	0.50	0.00	1.56	10.25	0.70	0.50	0.00	0.00	0.50	0.00	0.00	0.00	0.00	0.00	0.68	17.26			
DATA	DATA LAYER	0.30	0.00	0.40	0.00	0.00	2.00	0.00	0.00	0.00	0.00	1.50	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.66	4.86			
PIP	PIPELINES	2.82	0.50	0.00	0.00	1.50	0.00	0.00	0.00	0.50	1.80	2.25	0.00	0.20	0.00	0.00	0.00	0.60	0.00	0.00	0.00	2.10	2.35	14.62			
DELIV	DATA DELIVERY & TIERED DATA MODEL	1.00	0.00	1.35	0.50	0.00	0.00	0.20	0.00	0.00	0.00	0.00	0.25	0.00	0.00	0.00	0.00	0.00	3.78	0.00	0.00	0.00	0.89	7.97			
PROT	PROTOTYPING	0.10	0.25	0.00	2.00	0.00	0.00	1.10	1.00	0.00	0.00	0.00	0.85	0.00	0.00	0.60	2.50	1.40	0.00	0.00	0.00	0.00	0.83	10.63			
LMC	LOCAL MONITORING & CONTROL	0.00	0.00	0.00	0.00	0.00	0.50	0.00	0.00	0.00	0.00	1.00	0.00	1.20	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	2.70			
LINFA	Local Infrastructure	0.00	0.00	0.00	0.00	0.00	0.00	0.20	0.00	0.00	0.00	0.00	0.00	0.30	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.50			
		5.30	1.25	1.75	2.50	1.50	7.25	2.00	2.00	0.55	3.36	15.70	1.80	5.40	0.00	2.70	8.25	2.00	3.78	0.00	0.00	2.25	5.46	74.80			
		7.1%	1.7%	2.3%	3.3%	2.0%	9.7%	2.7%	2.7%	0.7%	4.5%	21.0%	2.4%	7.2%	0.0%	3.6%	11.0%	2.7%	5.1%	0.0%	0.0%	3.0%	7.3%				

A Stage-2 resources available review was carried out during May/June 2015. The results are shown above.

Requests to task leads for information about actual (visible) effort demonstrated per task has not provided sufficient information to carry out a gap analysis. This is also true in relation to resource needs going forward in Stage-2.

February 2014 view




SDP - Estimated steady state resourcing (03FEB14)


Units in the table are in person-months

Leader	L4 Label	L5 Label	Est Req	Allocated	Cover	UCAM	UOXF	UMAN	STFC	ASTN	CANU	CADC	CANARIE	SKSA	ICRAR	UCL	STON	VUW	AUT	MPIFR	JLCH	UFRN	CHPC	IVEC	BSC	PTC	CSIRO	UPRC	
Alexander	MGT		108	126	1.2	54																							
Ratcliffe	SE		209	199	1	18				27				36													14		10
Nikolic	ARCH		89	110	1.2	36	11			27																			
	Nikolic	ARCH			###																								
	Pickles	SWE	108	140	1.3	36			72																0	14	18		
	Zhao	EE	36	18	0.5																								18
	Calleja	OPS	24	51	2.1					30																			
Broekema	COMP		251	386	1.5	24	11			24				72	27							36		16	16	27			74
	Zhu	SCHED	24	36	1.5																								36
Wicenec	DATA		252	281	1.1		18			24		72											36						59
Nikolic	PIP				###																								
	Boonstra	INP	54	38	0.7					9				18									11						
	Nijboer	CAL	98	124	1.3	36	20			45	5																		
	Scaife	IMG	98	177	1.8					27	5			18															
	Stappers	NIP	100	103	1		36	42													25								0
	Salvini	CAS	36	58	1.6		18			24																			
	Johnston-Hollitt	SCI	120	100	0.8					15	0		0				13		72				0		16				
	Simmonds	UI	20	20	1								20																
Simmonds	DELIV		195	188	1		33			21	19.5	18	78	18															
Alexander	PROT				###																								
	Calleja	OAL	118	102	0.9		27	15		21													18	12		9			
	Simmonds	TDDL	126	138	1.1		9	15				19.5	78											16					
	Homberg/Fieseler	ISP	156	130	0.8		9		36														18	16	9				
	Bolton	SS	78	113	1.4	18	15							18		14						36	12						
Ratcliffe	LMC		75	92	1.2									54	18														20
Horrell	LINFA		39	42	1.1					6				18								18							
Totals			2414	2771	1.1	288	192	42	108	300	67	90	176	252	261	48	54	90	121	25	72	11	126	72	50	55	54	217	
as % project						10.4%	6.9%	1.5%	3.9%	10.8%	2.4%	3.2%	6.4%	9.1%	9.4%	1.7%	1.9%	3.2%	4.4%	0.9%	2.6%	0.4%	4.5%	2.6%	1.8%	2.0%	1.9%	7.8%	

SKA in Africa




African VLBI Network



Funded by DIRCO
(African Renaissance
Fund) and the DST



Ghana




Kenya

Big Data Africa

Mail & Guardian
AFRICA'S BEST READ

SKA takes the long view on big data

Sarah Wild



ASTRON & IBM Center for Exascale Technology
Drenthe, Netherlands



Big Data Africa

Activities

Funding
Networks
Workshops
Staff exchanges
Training
Resource pooling
Outreach, etc.

Research Themes (potential)

Astronomy
Smart Cities
Resource Management
Finance/commerce
Bioinformatics, etc.

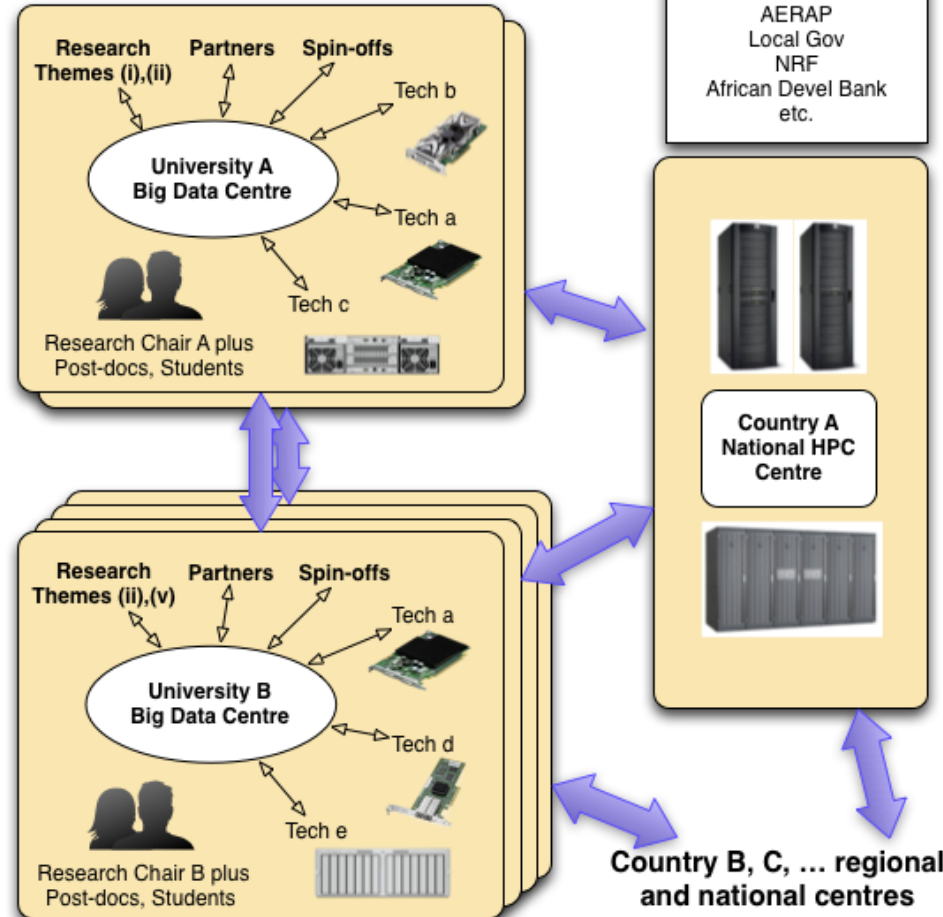
Machine Learning
Visualization
Storage
Streaming
Accelerators
Green computing, etc

Industry Partners (potential)

IBM
Intel
NVIDIA
Software AG
CISCO
Oracle
Cray
Huawei
etc.

Funding Partners (potential)

Industry
IDC
AERAP
Local Gov
NRF
African Devel Bank
etc.





Candidate compute architectures to address computational kernels and imaging pipelines:

- Many-core accelerators (e.g. GPGPU, FPGA and Xeon Phi)
- Arithmetic Processing Units comprising CPU and GPU in one package
- Low-power SoC technologies (e.g. ARM, Atom)

Storage Solutions addressing pseudo real-time buffering of the visibility buffer and the archive:

- Enterprise-level vs Commodity Disks (e.g. SAS vs. SATA)
- DRAM and Non-volatile (NVRAM) storage
- Parallel file systems (e.g. Lustre)
- Object-based storage (e.g. SWIFT, CEPH)

High performance networks addressing bulk-data transport and potential low-latency interconnect

- Infiniband and other “proprietary” networking
- High Speed Ethernet
- Software Defined Networks

Scheduling

Operations – service levels, system maintenance process, lifecycle management

OAL Development Plan Tasks

T-3: Operations Management

T-12: RAM analysis

T-28: Pseudo-real-time buffering of the visibility buffer and archival (hardware, file systems, object storage)

T-29: High Speed Ethernet and Software Defined Networks.

T-30: Energy Efficient Computing alternatives

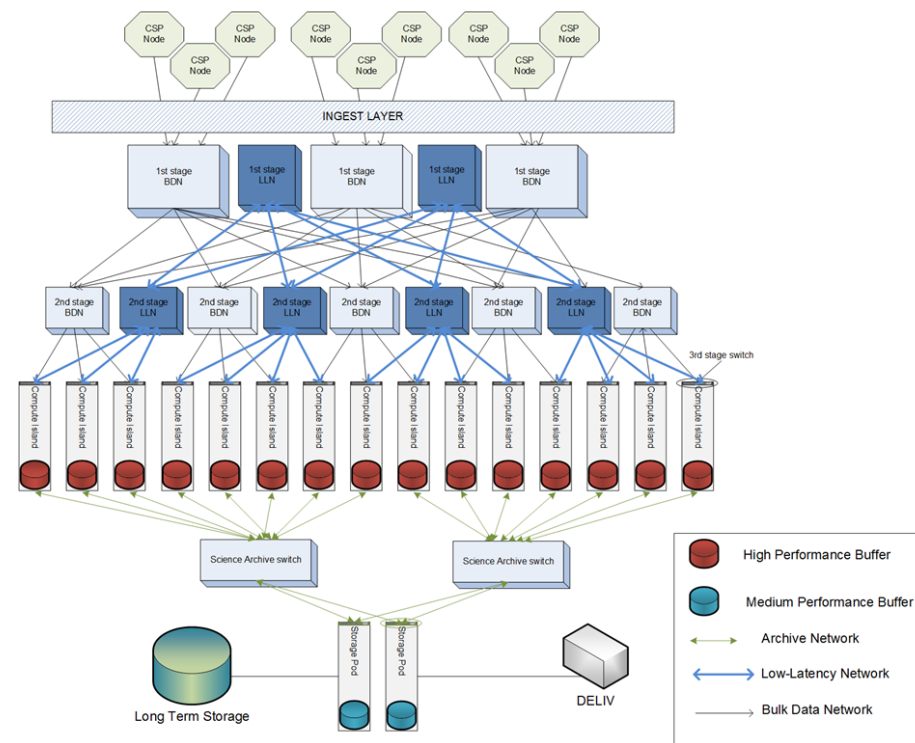
T-33: OpenStack and Containerization

T-34: LMC workflow management and scheduling

T-35: Assessment of Exascale system software alternatives

T-38: System Simulations.

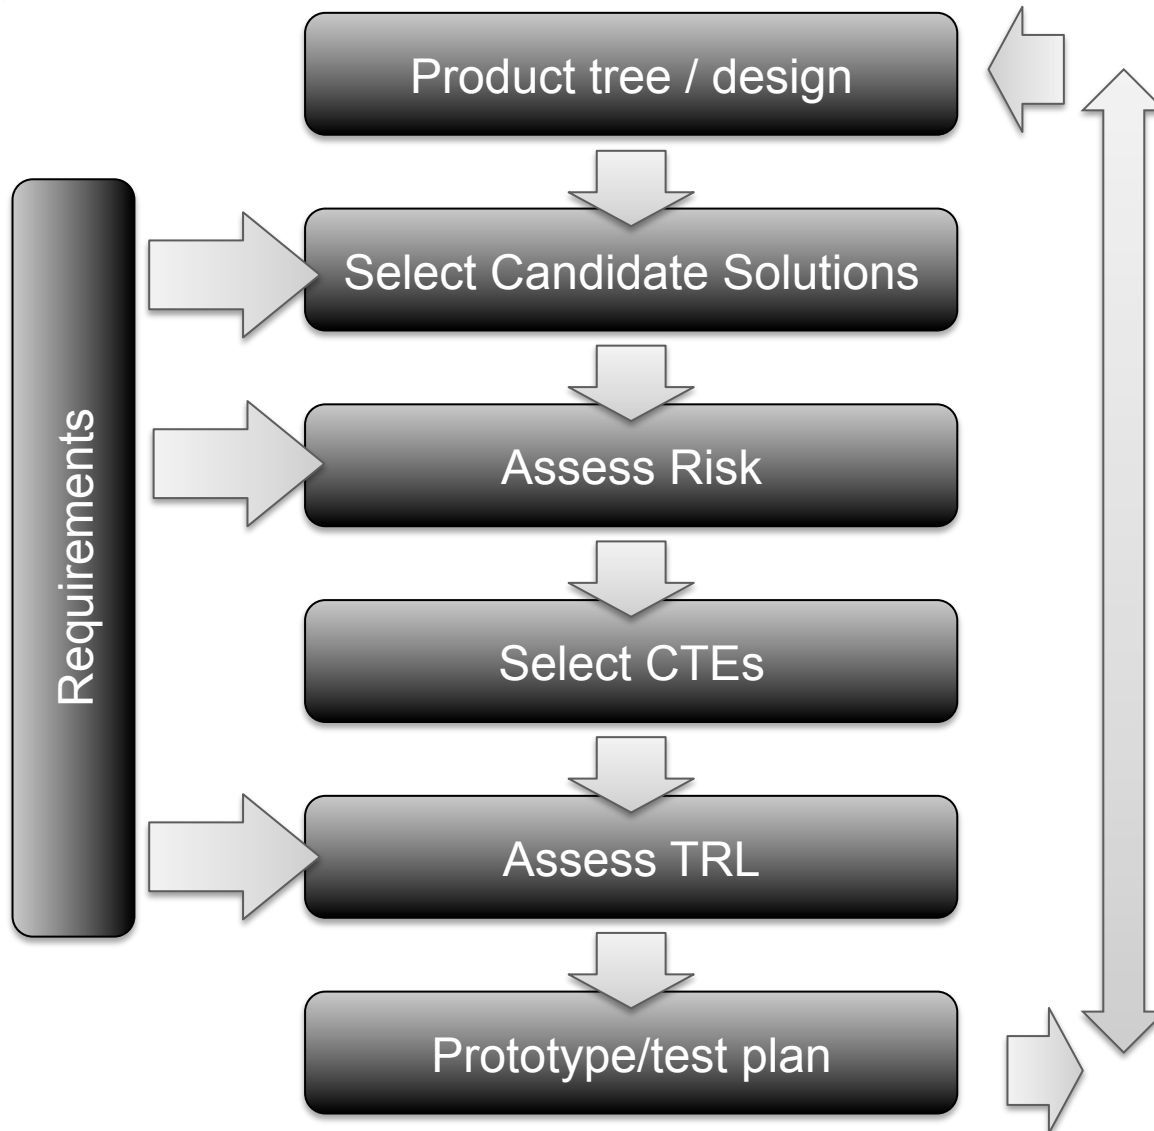
T-56: Candidate compute accelerators





- Initial Focus on key technology pinch-points – Processor, Storage, Networking and Data Flow
 - Many core/Accelerator model is seen as the most viable route – focused on algorithm development for Phi, GPGPU – APU/DSP work to follow (mostly x86, Power and ARM)
 - NVRAM – Initial work on CASA using SSD (2-3x over shared storage) and should be extended to track NvRAM technology
 - Networking – Tracking High-Arity Networking silicon and exploring QoS for combined networking
 - DataFlow

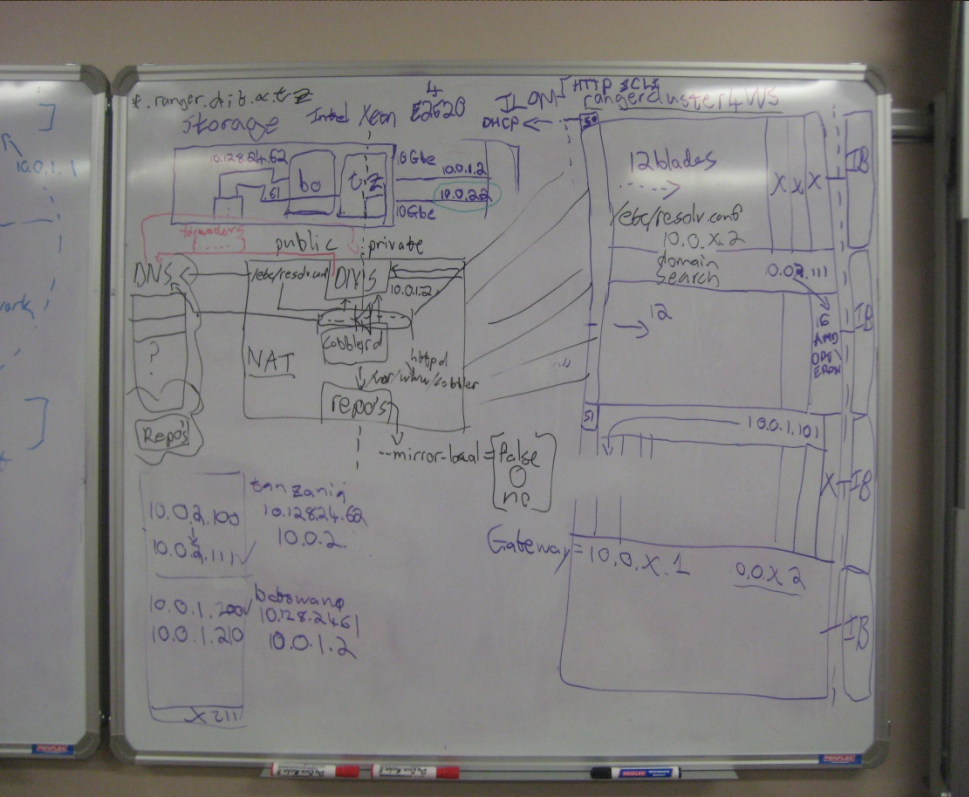
Process Overview





University of Botswana Computer Science Department

Building HPC Infrastructure & Partnerships for Research







System Information

- INSPUR Cluster – Intel X86
- 120 Nodes, 1700 Cores and 36 Tflops
- 2 GPU Nodes and 2MIC Nodes.

Applications

WRF and MM5 for Climate and Weather (Meteorological Services Department).
BLAST, jModelTest, PhyML and FigTree for Biosciences (African Institute of Biomedical Science & Technology (AIBST)).
VASP and Gaussian for Academia (University of Zimbabwe).



Collaboration efforts



1. 2013 - date , Participation in SADC Regional Collaborative Framework for High Performance Computing and Big Data
2. May 2014 UB Computer Science has received HPC equipment from Texas Advanced Computing Centre (TACC) to serves a basis for HPC service project Facilitated by CHPC and Dr Dan Stanzione and TACC
3. July 2014 UB Computer Science 4 person team hosted by CHPC for a 1 week workshop in Capetown around setting up TACC hardware - Nick Thorne
4. 2014/2015 UB Budget - DVCFA approved budget to buy supporting equipment operationalise the TACC Hardware based HPC.
5. **July - September 2015 - Setting up exploratory UB HPC Service**
6. July 2015 CHPC Post Graduate Workshop in Johannesburg, and invitation extended to Botswana by CHPC - Two Batswana students to attend
7. **29th July 2015 UB to host Newton Fund, the SKA and HPC delegation for discussions on HPC and interferometer sites**
8. **CHPC and UB Computer Science to schedule a Botswana HPC training workshop before the end of this year**
9. **November 2015 - Two or Three UB Computer Science Team invited to TACC and Texas for a technical workshop on received HPC hardware and for Supercomputing (SC-15) Conference - Facilitated by CHPC, SADC and TACC**

Proposed Roll-out Plan



Item	Time-line	Responsible Party
HPC roll-out for Phase 1 (Botswana, Zambia, Namibia and Mauritius)	1 st April 2015 – 30 th March 2016	CHPC
Training Workshops	June 2015	CHPC and SKA South Africa
HPC roll-out for Phase 2 (Ghana, Madagascar, Kenya and Mozambique)	1 st April 2016 to March 2017	CHPC
Training workshop	June 2016	CHPC and SKA South Africa
Deployment of Astronomy Software	Will be done as infrastructure is deployed	SKA South Africa and international partners.
Up-scaling of the CHPC Research Cloud	1 st April 2015	CHPC and SKA South Africa





- SADC HPC Framework
- CHPC Annual Meeting (small scale of ISC and SC)
- Student Development (SCC)
 - Cooperation with DELL with multi-year
 - Cooperation with Mellanox on switches
 - Success in identifying good skills in the country
- Collaboration with TACC proved successful and promising to strengthen developments.
- SKA Consortium gives opportunity to evaluate diverse architecture at different scales.