



Architect of an Open World™

## The RWTH Aachen Supercomputer Architecture and Design

**Dr. Thomas Warschko, Technical Director Extreme Computing**

**PPCES – March 11th 2013**



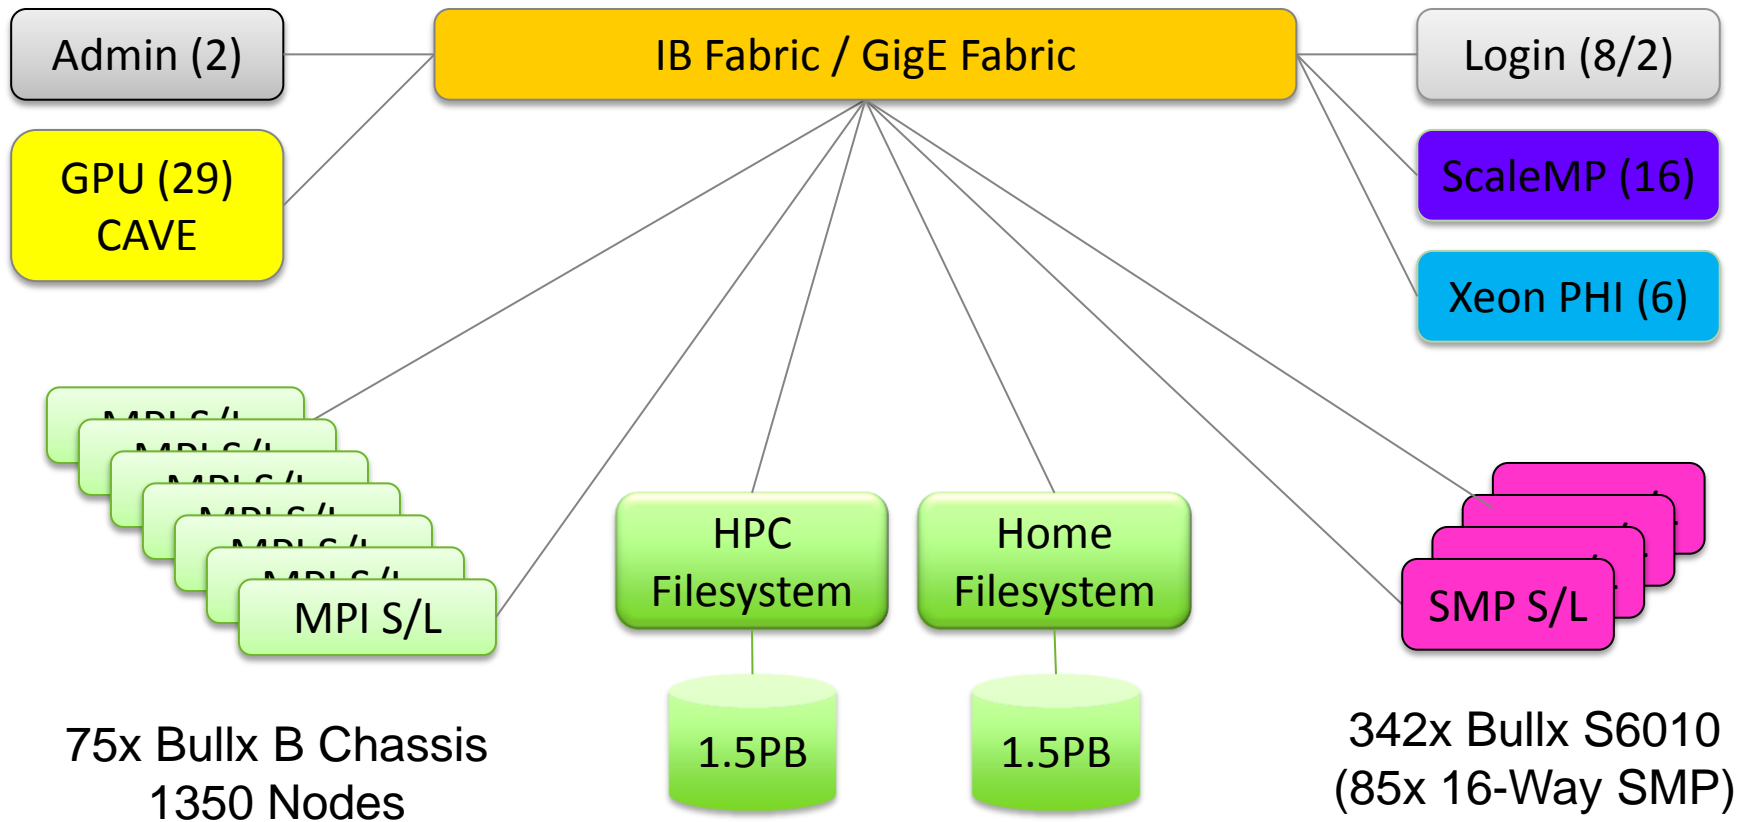
# The RWTH Aachen Supercomputer at a glance

Architect of an Open World™



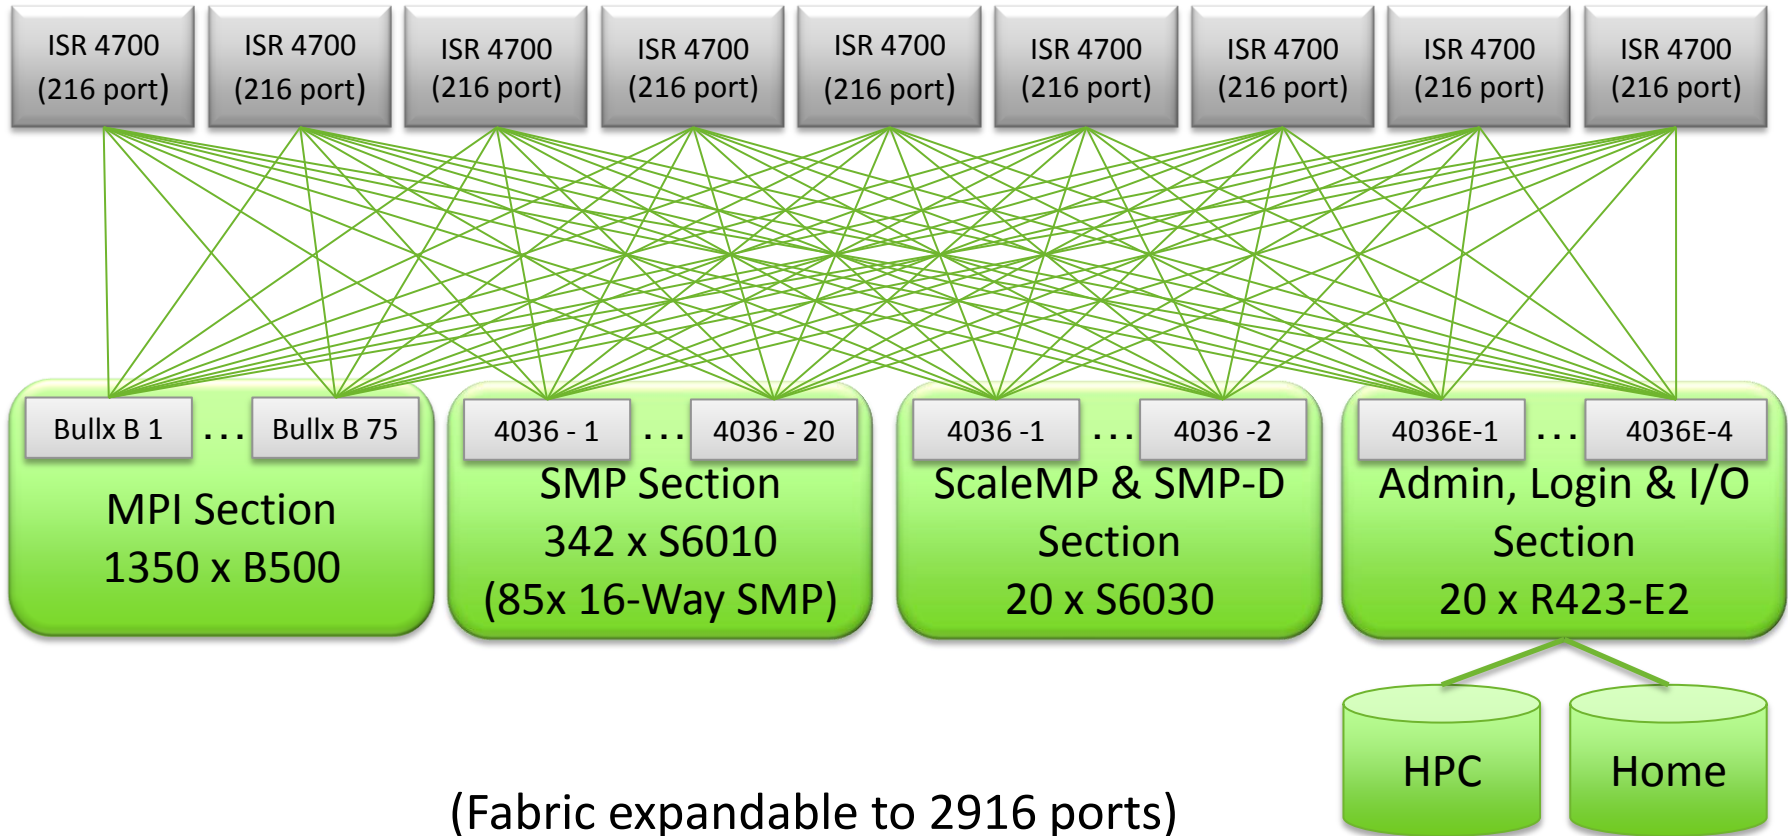
292	Tflop/s Peak
220	Tflop/s Linpack
1 350	bullx blades
362	bullx S modules
16 200	Intel Westmere cores
11 500	Intel Nehalem-EX cores
1 500	TB HPC storage
1 500	TB Home storage
	QDR InfiniBand interconnect
29	GPU Nodes with 2x Nvidia Q6000 +Gsync
6	Accelerator Nodes with 2x Xeon-PHI 5110P

# RWTH Aachen – System Architecture



# Infiniband Fabric – Full Fat Tree

9x 216port TopLevel Switches - 1944 ports total




# The System in Real



# MPI Section: Bullx B500 Blade Series

- 75x Bullx Blade Chassis
- 1350x Bullx B500 Compute Node
  - 1098 Nodes with 24 GB Memory
  - 252 Nodes with 96 GB Memory
- Intel Westmere X5675 CPU (3,06GHz)
- Local SSD Disk (60GB)
- QDR Infiniband



Bullx B 1



...

Bullx B 75

MPI Section  
1350 x B500

# SMP Section: Bull S6010 Supernode

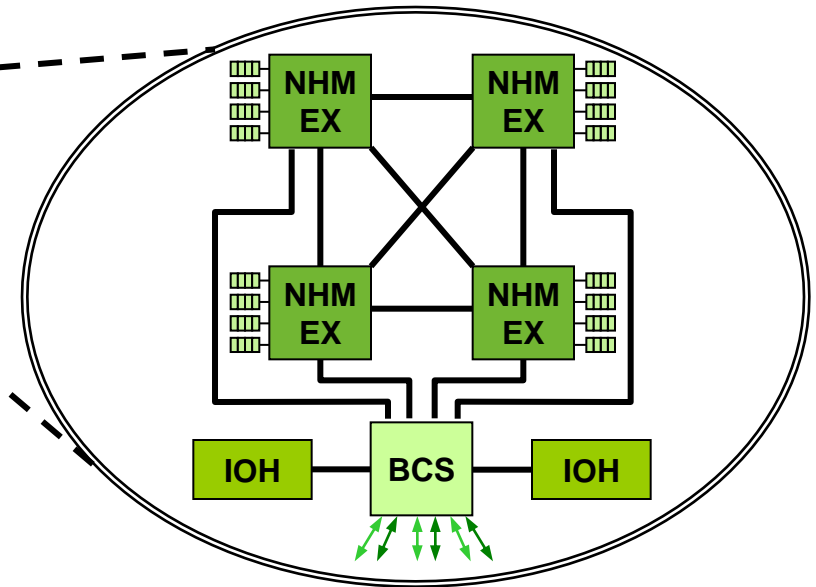

- 342x Bullx S6010 Compute Node
  - 270 Modules with 64 GB Memory (67x 16way – 256 GB Memory)
  - 72 Modules with 256 GB Memory (18x 16way – 1024 GB Memory)
- Intel Nehalem-EX X7550 CPU (2 Ghz)
- Local Sata Disk (500GB)
- QDR Infiniband



4036 - 1 ... 4036 - 20

SMP Section  
342 x S6010  
85x 16-Way SMP

# Bullx S60x0 – CC-NUMA server



Node maximum configuration :


- 4 modules
- 16 sockets
- 128 cores (Nehalem-EX)
- 128 memory slots (2TB)

Large nodes :

- Large shared memory (pre/post-processing)
- Many more cores (SMP)
- Fewer nodes
- Simpler system administration
- Multi Level Parallelism (MPI/OpenMP)

# ScaleMP Section: Bullx S6030 Supernode

- 16x Bullx S6030 Compute Node
  - 256 GB Memory
  - Intel Nehalem-EX X7550 CPU (2 Ghz)
  - 3x 300GB SAS Disk
  - QDR Infiniband
  - 10GigE
- Single vSMP system



4036 -1 ... 4036 -2

ScaleMP Section  
16 x S6030

## Service and Admin Section: Bullx R423-E2

- **8x MPI-D: R423-E2**
  - Intel Westmere X5675, 96GB Memory  
5x300 GB SAS Disk, QDR-IB
- **4x SMP-D: S6030**
  - Intel Nehalem-EX X7550, 128GB Memory  
3x300GB SAS Disk, QDR-IB
- **2x Admin: R423-E2**
  - Intel Westmere X5675, 48GB Memory  
4x600 GB SAS Disk, QDR-IB
- **8x OSS: R423-E2**
  - Intel Westmere X5675, 48GB Memory  
2x300 GB SAS Disk, 3x QDR-IB
- **2x MDS: R423-E2**
  - Intel Westmere X5675, 48GB Memory  
2x300 GB SAS Disk, QDR-IB, 2x FC8



# HPC & Home Filesystem: DDN & NetApp

## ■ HPC Filesystem / DDN:

- 2x DDN SFA10k
- 940x 2TB SATA Disk
- 1x DDN 6620 (as MDS)
- 40x 600GB SAS Disk 15krpm

## ■ Home Filesystem / NetApp:

- 6x FAS 6280 NAS Filer
- 54x DiskShelves
- 1296x 2TB SATA Disk



# System Summary

Section	Component / CPU	#Nodes/#Modules	TFLOP
MPI-S	B500 / X5675	1098	161
MPI-L	B500 / X5675	252	37
MPI-D	R423-E2 / X5675	8	1
SMP-S	S6010 / X7550	67/270	69
SMP-L	S6010 / X7550	18/72	18
SMP-D	S6030 / X7550	4	1
ScaleMP	S6030 / X7550	16	4
		$\Sigma = 1712$	$\Sigma = 292$ TFlop
Home-Filesystem	NetApp / FAS6280	1.500 TB	
HPC-Filesystem	DDN / SFA10k	1.500 TB	
Infiniband	Voltaire	Full Fat-Tree	2000 ports

## Top 150 Entries - Germany (Top500 - Nov 2012)

Rank	Site	Cores	Rmax	Rpeak
5	FZ Jülich (JUQUEEN)	393.216	4.141,2	5.033,2
6	LRZ (SuperMUC)	147.456	2.897,0	3.185,1
27	HWW (HERMIT)	113.427	831,4	1.043,9
84	University of Frankfurt	15.120	285,2	469,7
89	FZ Jülich (JUROPA)	26.304	274,8	308,3
110	Mainz (MOGON)	33.920	225,6	284,9
<b>111</b>	<b>RWTH Aachen</b>	<b>25.488</b>	<b>219,8</b>	<b>270,5</b>
128	Max-Planck	9.904	187,4	206,0



## Top 500 Entry - RWTH (2011 - 2013)

List	Rank (World)	Rank (Germany)
June 2011	32	4
Nov 2011	47	5
June 2012	78	7
Nov 2012	111	7



Architect of an Open World™

Thank you





# bullx

instruments for innovation

