


ScaleMP

Dirk Schmidl

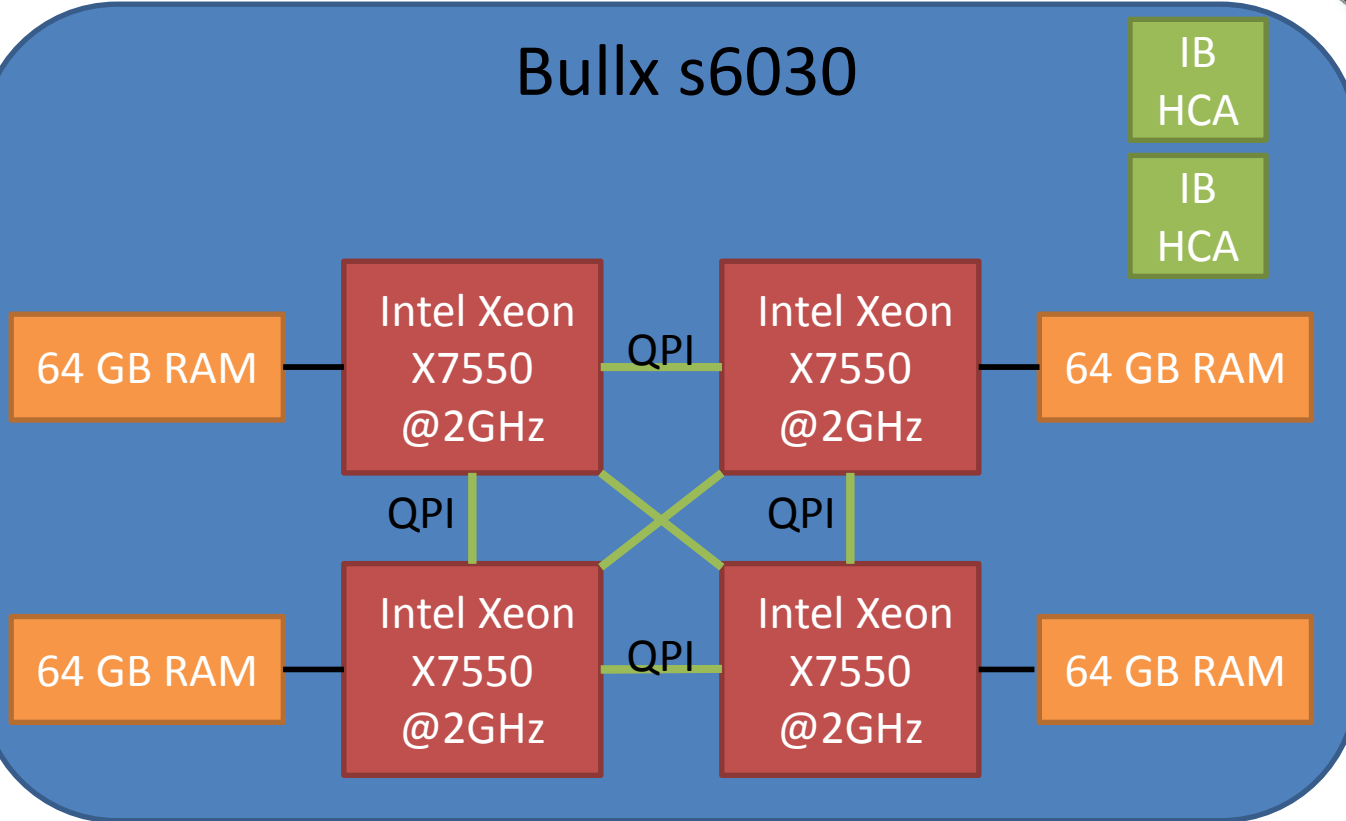
schmidl@rz.rwth-aachen.de

PPCES / 21.03.2012 / Aachen

Bullx s6030 board



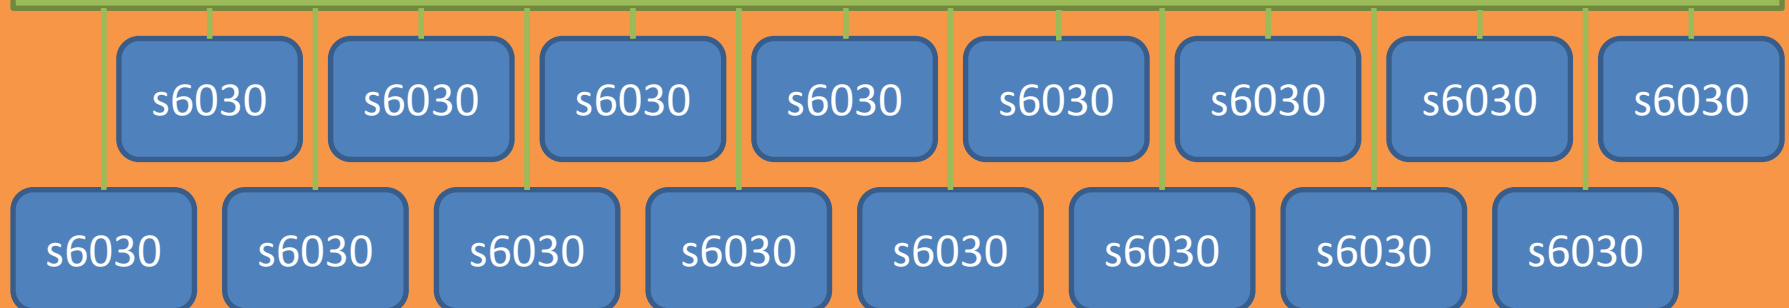
Bullx s6030



vSMP Foundation:

- creates a Single System image on top of a standard x86 based cluster
- parts of the memory are used for prefetching and caching by vSMP software
- remote access, prefetching and caching works on a 4k basis
- infiniband network is required internally for data traffic

4xQDR Infiniband



Recourse	ScaleMP	Bull BCS
CPUs	64	16
Cores	512	128
Rate	2 GHz	2 GHz
Memory	~ 4 TB	256 GB
Total L3 Cache	1.2 GB	288 MB
Peak Performance	4 TFLOP/s	1 TFLOP/s
Single System Image provided by	vSMP Software	Bull Coherence Switch (Hardware)

- ▶ Standard Benchmark to measure memory performance.
- ▶ Version is parallelized with OpenMP.

Measures Memory bandwidth for:

y=x (copy)

y=s*x (scale)

y=x+z (add)

y=x+s*z (triad)

```
#pragma omp parallel for
for (j=0; j<N; j++)
    b[j] = scalar*c[j];
```

for double vectors x,y,z and scalar double value s

Function	Rate (MB/s)	Avg time	Min time	Max time
Copy:	33237.0185	0.0050	0.0048	0.0055
Scale:	33304.6471	0.0049	0.0048	0.0059
Add:	35456.0586	0.0070	0.0068	0.0073
Triad:	36030.9600	0.0069	0.0067	0.0072

Stream (11 GB Memory Footprint)

Initialization:

```
for (j=0; j<N; j++)  
{  
    a[j] = 1.0;  
    b[j] = 2.0;  
    c[j] = 0.0;  
}
```

Triad:

```
#pragma omp parallel for  
for (j=0; j<N; j++)  
{  
    b[j] = scalar*c[j];  
}
```

\$ OMP_NUM_THREADS=128 ./stream

	no Pinning/ serial Init	Pinning/ serial Init	Pinning/ parallel Init
BullBCS 128T	9 GB/s		
ScaleMP 128T	18 GB/s		
ScaleMP 512T	120 GB/s		

Stream (11 GB Memory Footprint)

Initialization:

```
for (j=0; j<N; j++)  
{  
    a[j] = 1.0;  
    b[j] = 2.0;  
    c[j] = 0.0;  
}
```

Triad:

```
#pragma omp parallel for  
for (j=0; j<N; j++)  
{  
    b[j] = scalar*c[j];  
}
```

\$ OMP_PROC_BIND=true OMP_NUM_THREADS=128 ./stream

	no Pinning/ serial Init	Pinning/ serial Init	Pinning/ parallel Init
BullBCS 128T	9 GB/s	7 GB/s	
ScaleMP 128T	18 GB/s	47 GB/s	
ScaleMP 512T	120 GB/s	149 GB/s	

Stream (11 GB Memory Footprint)

Initialization:

```
#pragma omp parallel for
for (j=0; j<N; j++)
{
    a[j] = 1.0;
    b[j] = 2.0;
    c[j] = 0.0;
}
```

Triad:

```
#pragma omp parallel for
for (j=0; j<N; j++)
{
    b[j] = scalar*c[j];
}
```

```
$ OMP_PROC_BIND=true OMP_NUM_THREADS=128 ./stream
```

	no Pinning/ serial Init	Pinning/ serial Init	Pinning/ parallel Init
BullBCS 128T	9 GB/s	7 GB/s	236 GB/s
ScaleMP 128T	18 GB/s	47 GB/s	235 GB/s
ScaleMP 512T	120 GB/s	149 GB/s	678 GB/s