

IBM Spectrum LSF Mehr als nur ein Scheduler


ZKI-Arbeitskreis Supercomputing
Herbsttreffen 2021
8.10.2021



Christof Westhues
IBM Spectrum Computing



IBM Spectrum Computing
products are used by 23 of the
30 largest commercial
enterprises in the world
including ...



**Numerous
Genomics,
Bioinformatics,
and Children's
Hospitals**



**Top 3
Cancer Centers in
the US**



**10 of top 12
Automotive
Companies**



**23 of top 25
Electronic
Companies**



**7 of top 10
Aero & Defense
Companies**



IBM Spectrum Computing



High Performance Computing
Design / Simulation / Modeling

High Performance Analytics
Trade / Risk Analytics

'New-gen Workloads'
Hadoop, Spark, Containers

Workload
Aware
Scheduling

Shared
Resource
Management



IBM
Spectrum
LSF



IBM
Spectrum
Symphony




IBM
Spectrum
Conductor



IBM Spectrum Computing

Heterogeneous Servers & Storage




Flash



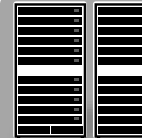
Disk




Tape




Power



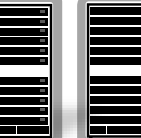
x86




Linux on z




Sparc



ARM



docker



VM



IBM Cloud
Microsoft Azure

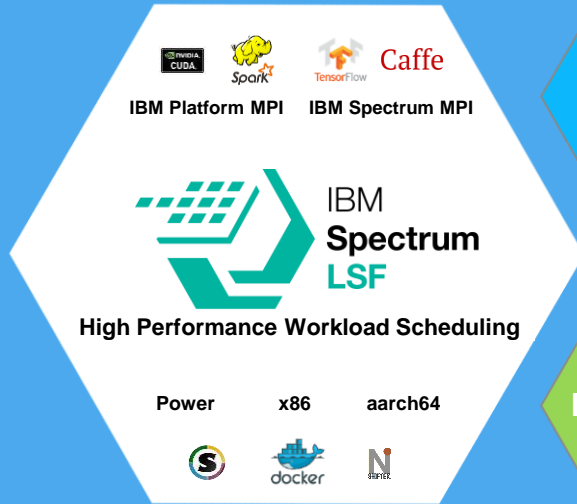
Hybrid Cloud Infrastructure

IBM Spectrum LSF Suites 10.2

Suite for Enterprise

Suite for HPC

Suite for Workgroups



Workgroup Suite components:

- Logos: NVIDIA CUDA, Spark, TensorFlow, Caffe
- Text: IBM Platform MPI, IBM Spectrum MPI
- Logo: IBM Spectrum LSF
- Text: High Performance Workload Scheduling
- Text: Power, x86, aarch64
- Logos: SUSE, docker, NIX

Application
Centric User
Portal

Lightweight
Reporting &
Dashboards

Cluster
Deployment and
Monitoring

Workflow
Automation

Intelligent
Data Staging

Hybrid Cloud
Auto-scaling

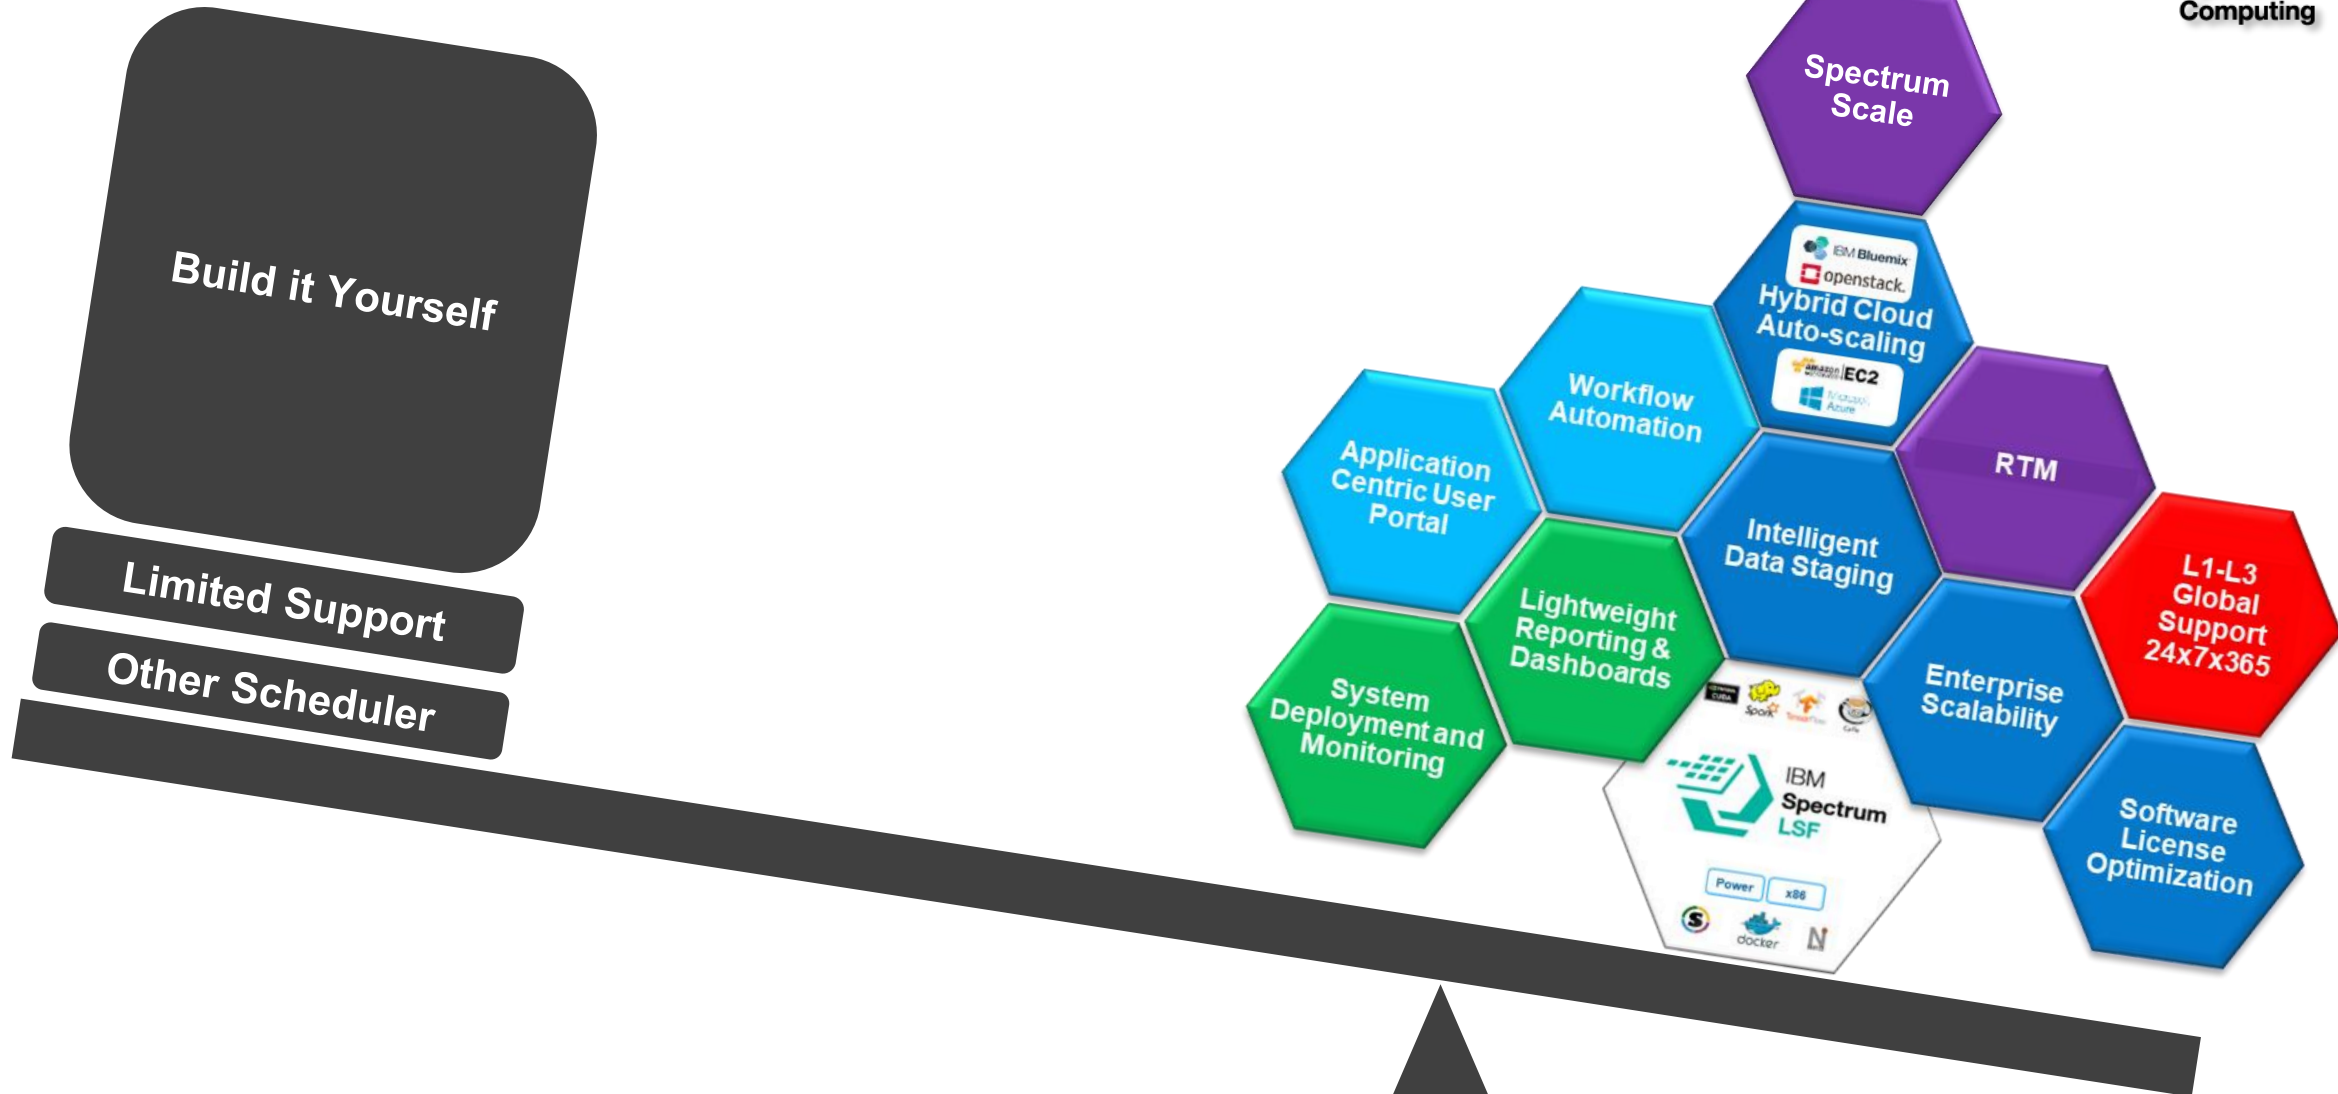
Software
License
Optimization

Enterprise
Scalability

IBM Spectrum Software - complete end to end solution



IBM
Spectrum
Computing



Spectrum LSF Over 25+ Years of Sustained Agility!

Over 2,000 *Customer Driven*
enhancements since 1992!

“A product imagined by customers and
built by IBM”



IBM Spectrum LSF – The HPC Swiss Army Knife



Key LSF Features Purpose Built for HPC

- ✓ Unmatched Scheduler Performance
- ✓ Cloud Bursting (Host Factory)
- ✓ Control Groups and Affinity
- ✓ Container Support (Docker, etc.)
- ✓ Compute Units & Job Packing
- ✓ Event/Relaxed Scheduling
- ✓ 9th Generation GPU Support*
- ✓ Spectrum LSF Application Center
- ✓ Spectrum LSF Data Manager
- ✓ Global Fairshare
- ✓ Multi Cluster
- ✓ Session Scheduling
- ✓ Block Scheduling
- ✓ Multi-Pass Resource Requirements
- ✓ Absolute Priority Scheduling
- ✓ Spectrum LSF RTM
- ✓ Guaranteed Resources
- ✓ Multi-threaded Dispatch
- ✓ Key Pending Reasons
- ✓ Start Time Prediction
- ✓ Dynamic Resource Scheduling
- ✓ Resource Preemption
- ✓ License Scheduling
- ✓ Resizable Jobs
- ✓ Backfill Scheduling
- ✓ Delegation of Administrative Authority
- ✓ Alternative Resource Requirements
- ✓ Application Profiles
- ✓ Large Array Support
- ✓ Plan Based Scheduling
- ✓ LSF Simulator

NVIDIA GPGPU Support

Enforcement of GPU allocations via Control Groups (cgroups)

Dynamic mode switching

Exclusive allocation and round robin shared mode allocation.

CPU-GPU Affinity

Boost Control

Power Management

Memory Reservation


Zero Config Installation

NVLink Affinity

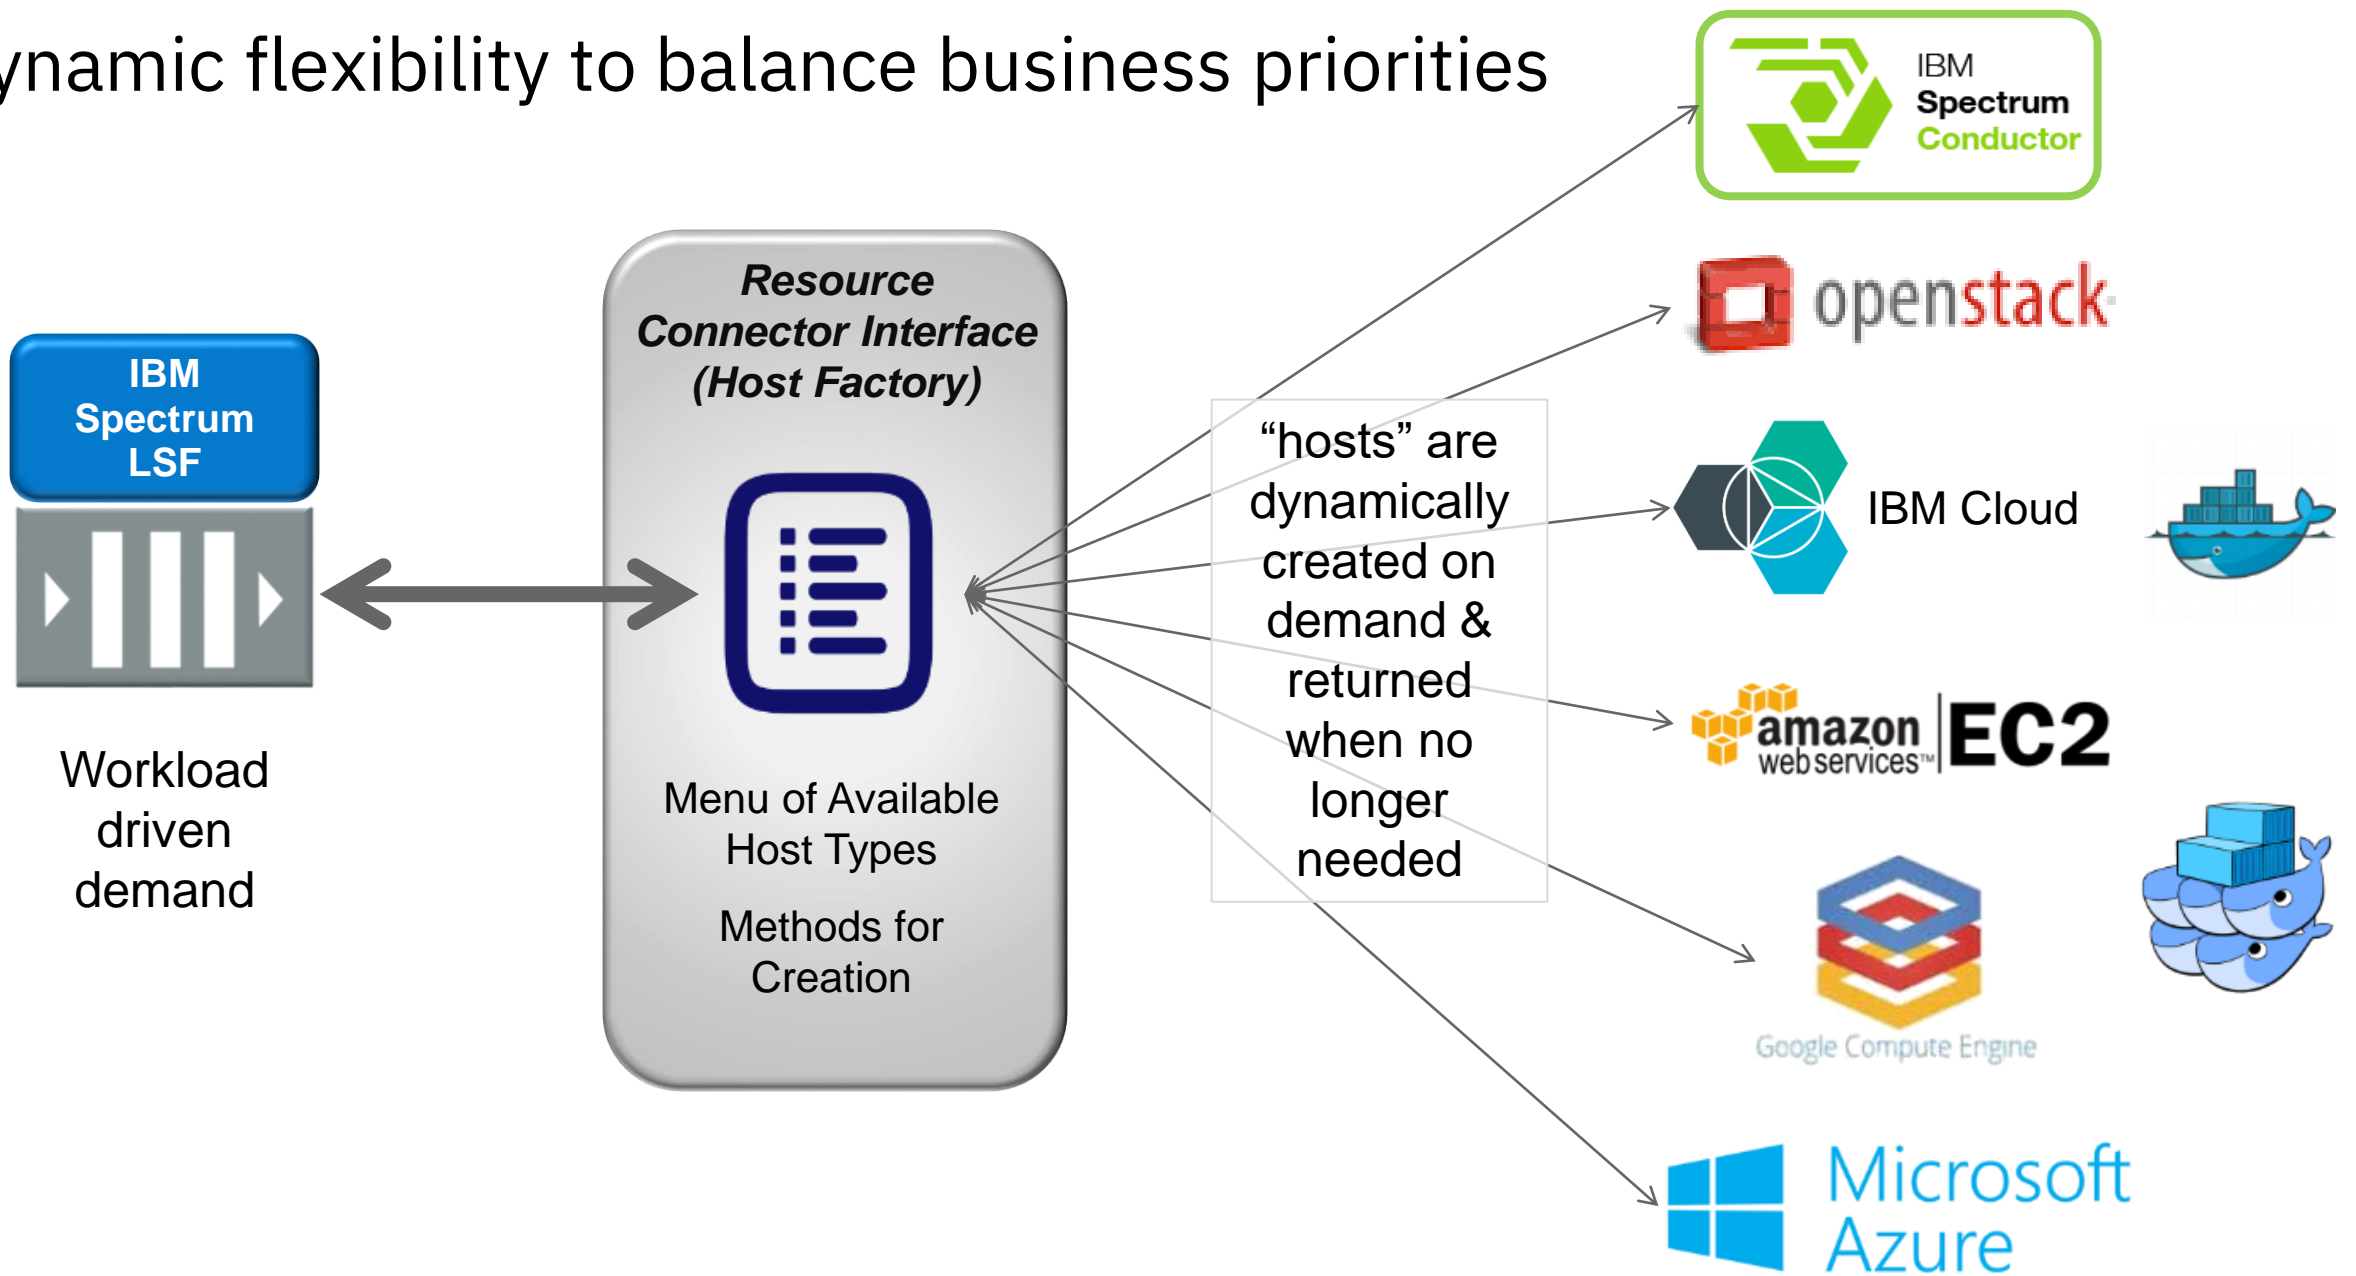
Multi-Process Server (MPS) Support

Pascal & DCGM Support

- ✓ GPU behavior monitoring
- ✓ GPU configuration management
- ✓ GPU policy oversight
- ✓ GPU health and diagnostics
- ✓ GPU accounting and process statistics





Dynamic flexibility to balance business priorities



Cloud Bursting with LSF





IBM
Spectrum
Computing




Job Submission

LSF MultiCluster forwards specific workloads to the cloud based on site defined policies.

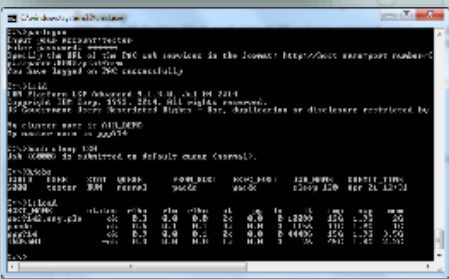



Single Connection between on-prem and Cloud





On Premise On Cloud

Cloud Bursting with LSF





Job Submission

LSF MultiCluster forwards specific workloads to the cloud based on site defined policies.



LSF's Resource Connector dynamically resizes the pool of hosts in the cloud based upon workload demands and policies.



On Premise On Cloud

Benefits of IBM Spectrum LSF container support:

Enable Docker jobs to run on IBM Spectrum LSF managed clusters

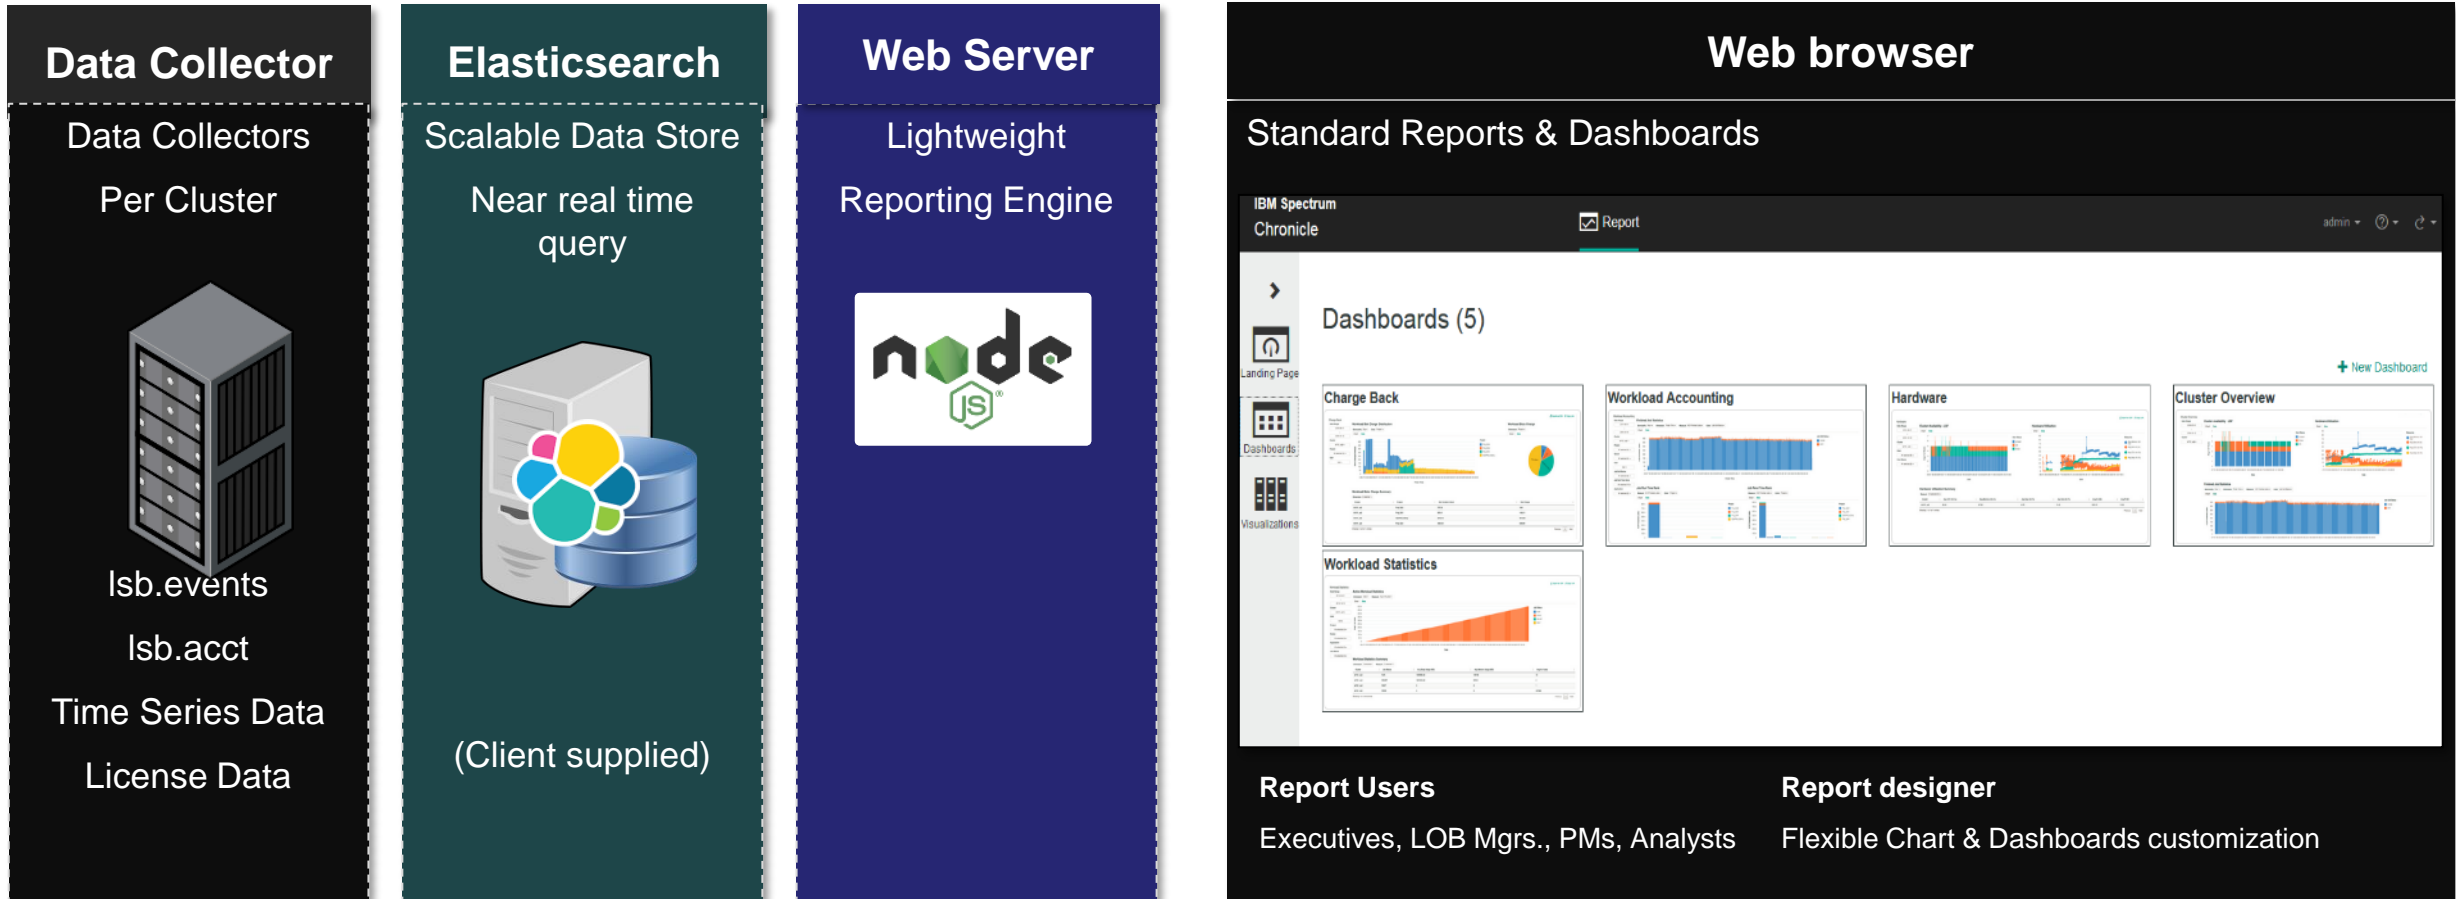
- Straightforward Docker enablement. Easy to install and easy to manage
- Allowing HPC users to run Docker jobs in batch mode in the same way as running other un-Dockerized jobs.
- All features in existing LSF V10.1 release are available to Docker jobs
- Addresses security concerns with Docker

Enabling system administrators to withhold root privileges from HPC user.

- On compute cluster: HPC users are no longer in Docker user group
- On development node(s) : HPC users remain in Docker user group to develop apps.
- Effectively prevented potential security and privacy violations from Docker users

Seamless integration with the rest of LSF family of software

IBM Spectrum LSF Explorer – Overview




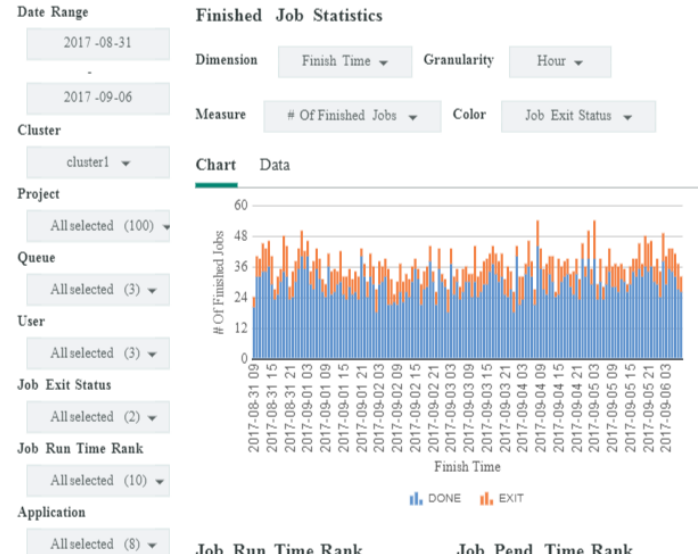
The architecture is aligned with other LSF solutions:

- LSF ELK stack for LSF cluster monitoring
- Elasticsearch-based LSF command utilities (bhist, bacct) to facilitate query and troubleshooting in second(s)


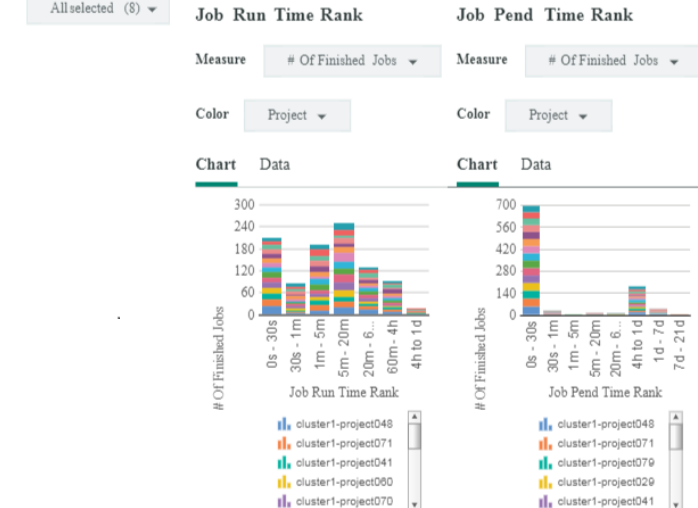
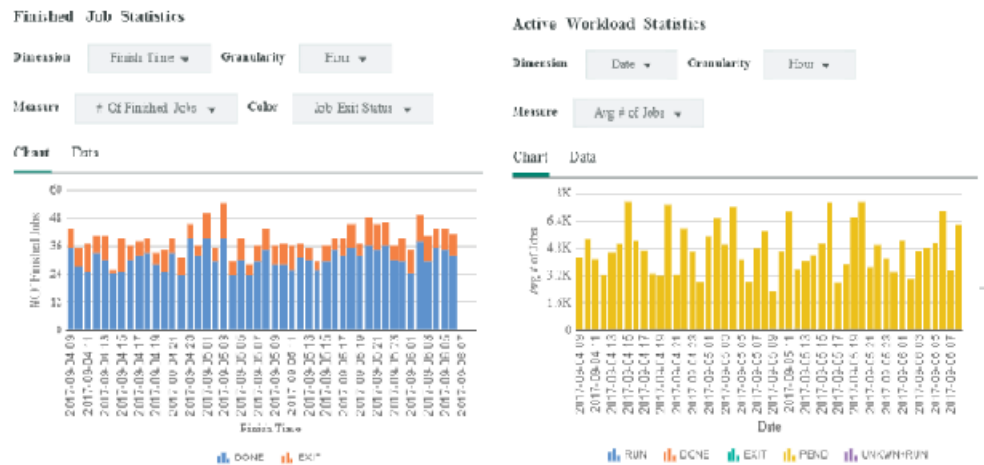
IBM Spectrum LSF Explorer – Sample charts



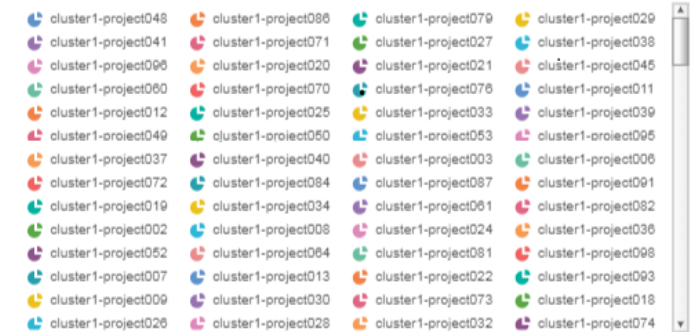
IBM
Spectrum
Computing



Job Run Time Rank



Workload Accounting Charts





Cluster Overview Charts


Workload Accounting Charts

Workload Slot Charge

Orchestrating Resources for Deep Learning (Reducing Jitter) Before LSF

Data scientist

qsub -n 10 ./train.py

Researcher

qsub -n 10 ./train.py



Orchestrating Resources for Deep Learning (Reducing Jitter) After LSF

Data scientist

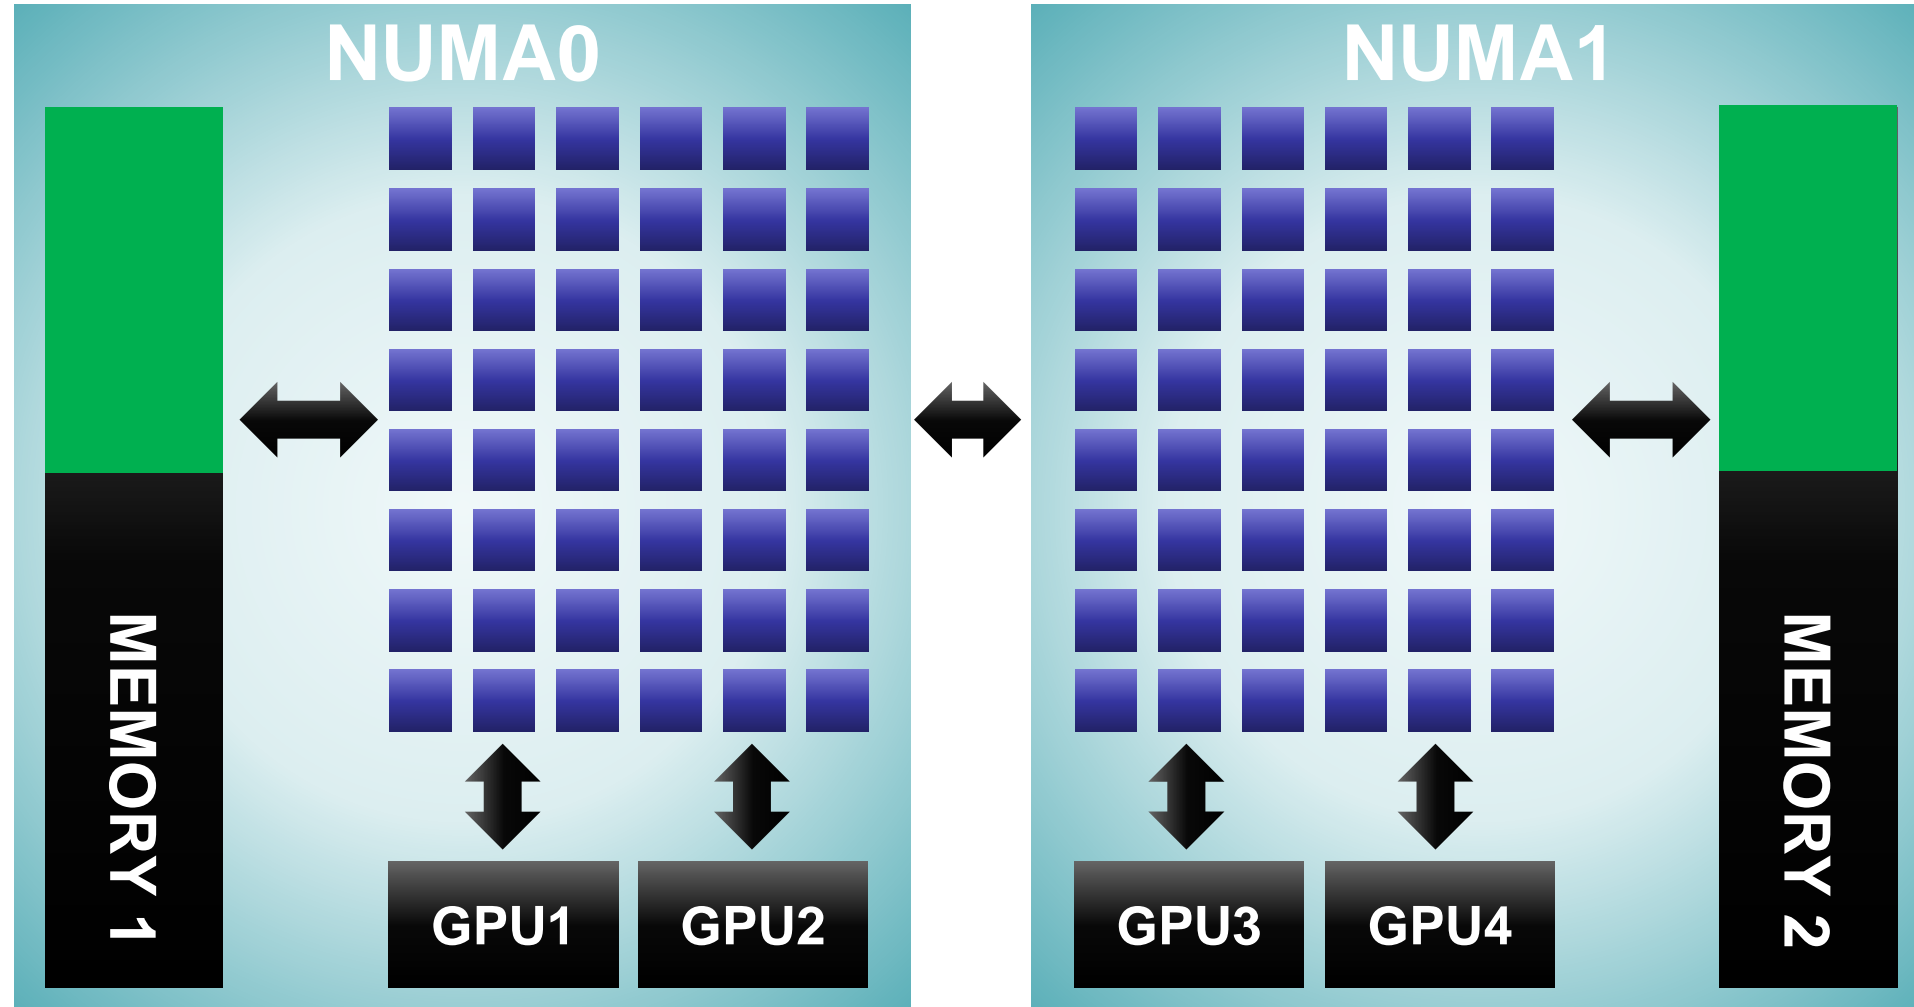


```
bsub -n 10 \  
-R "affinity[core(same=numa): \  
distribute=pack]" \  
-gpu "num=1:mode=exclusive" \  
./train.py
```

Researcher



```
bsub -n 10 \  
-R "affinity[core(same=numa): \  
distribute=pack]" \  
-gpu "num=1:mode=exclusive" \  
./train.py
```



Webinar am 21.Oktober 15 Uhr

SMB & IBM laden Sie ein zum Webinar
„IBM Spectrum LSF -
Mehr als nur ein Scheduler
Workload Management für HPC

Wir zeigen Ihnen wie Sie mit diesem professionellen Werkzeug das Maximale aus Ihrer HPC Investition herausholen und wie einfach sowohl User und Administratoren Workloads managen können. Teilnehmer erfahren außerdem wie sie das Werkzeug kostenlos ausprobieren können und wie wir bei den ersten Schritten begleiten.

Anmeldung:

<https://ibm.webex.com/ibm/onstage/g.php?MTID=ea56b2fcaacbb004cfba12000c0de6564>

Oder:

