




IBM Spectrum Conductor with Spark

Reference Architecture



Agenda

- Network Configuration
- Hardware Bill of Materials
- Software Bill of Materials

Recommended Network Configuration

Three Networks:

1. Provisioning & Management

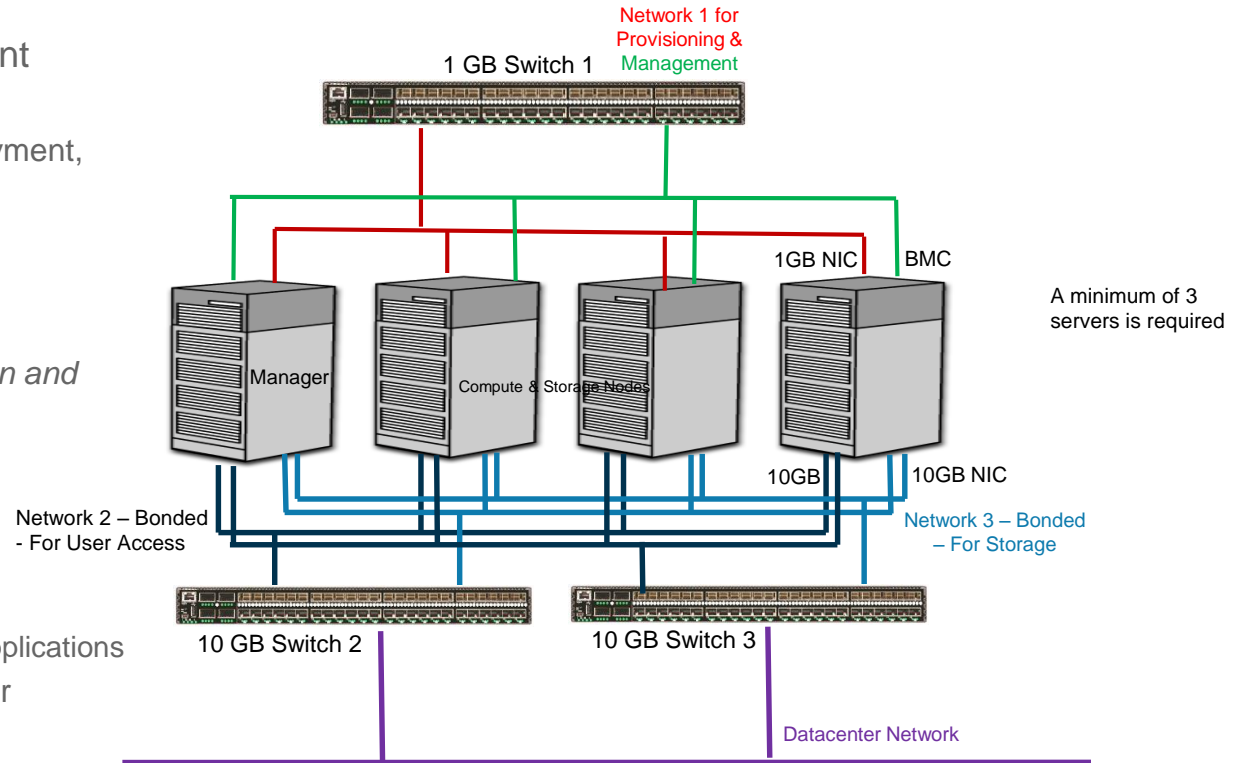
- Internal to appliance
- Used for appliance deployment, scale-out, monitoring and management

2. Storage

- *Spectrum Scale replication and access traffic*
- Bonded for reliability
- Jumbo frames

3. Access

- User access to the Spark applications
- Routable to the datacenter
- Bonded for reliability



Recommended Hardware Bill Of Materials

	Storage Rich ++	Compute Rich ++
CPU	Dual socket, 16+ core	Dual socket, 16+ core
Memory	256+ GB**	256+ GB**
Disks	2 - 500GB SSD SATA disks (RAID 1 for OS with 200GB partition for HA)	2 - 500GB SATA disks (RAID 1 for OS with 200GB partition for HA)
	2-4 – 1TB NVMe or SSDs for Spark Spillover	2-4 – 1TB NVMe or SSDs for Spark Spillover
	24+ - 6 TB SATA disks + RAID controller (For Spectrum Scale)*	0+ - 6 TB SATA disks + RAID controller (For Spectrum Scale)*
GPU (Optional)	Nvidia CUDA card	Nvidia CUDA card

++ A minimum of three servers is required. If any storage rich servers are to be used, then the first three servers must be storage rich. After that they may be a combination of storage rich or compute rich.

** Filling all the DIMM slots may provide better performance than having some slots open

* The sum of the virtual disks used for shared storage should not exceed 1500, or 18PB

Recommended Hardware Bill Of Materials

	Storage Rich	Compute Rich
Network <i>(all servers need to have the same NICs)</i>	1 - 1 GB NIC for management	1 - 1 GB NIC for management
	4 - 10GB NICs Bonded + VLAG 2 for Spark access 2 for Storage (Optional Jumbo Frames)	4 - 10GB NICs Bonded + VLAG 2 for Spark access 2 for Storage (Optional Jumbo Frames)
Network Switches	1 - 1GB switch for OS provisioning and BMC access	1 - 1GB switch for OS provisioning and BMC access
	2 - 10GB switches partitioned into Storage and Access VLANs.	2 - 10GB switches partitioned into Storage and Access VLANs.

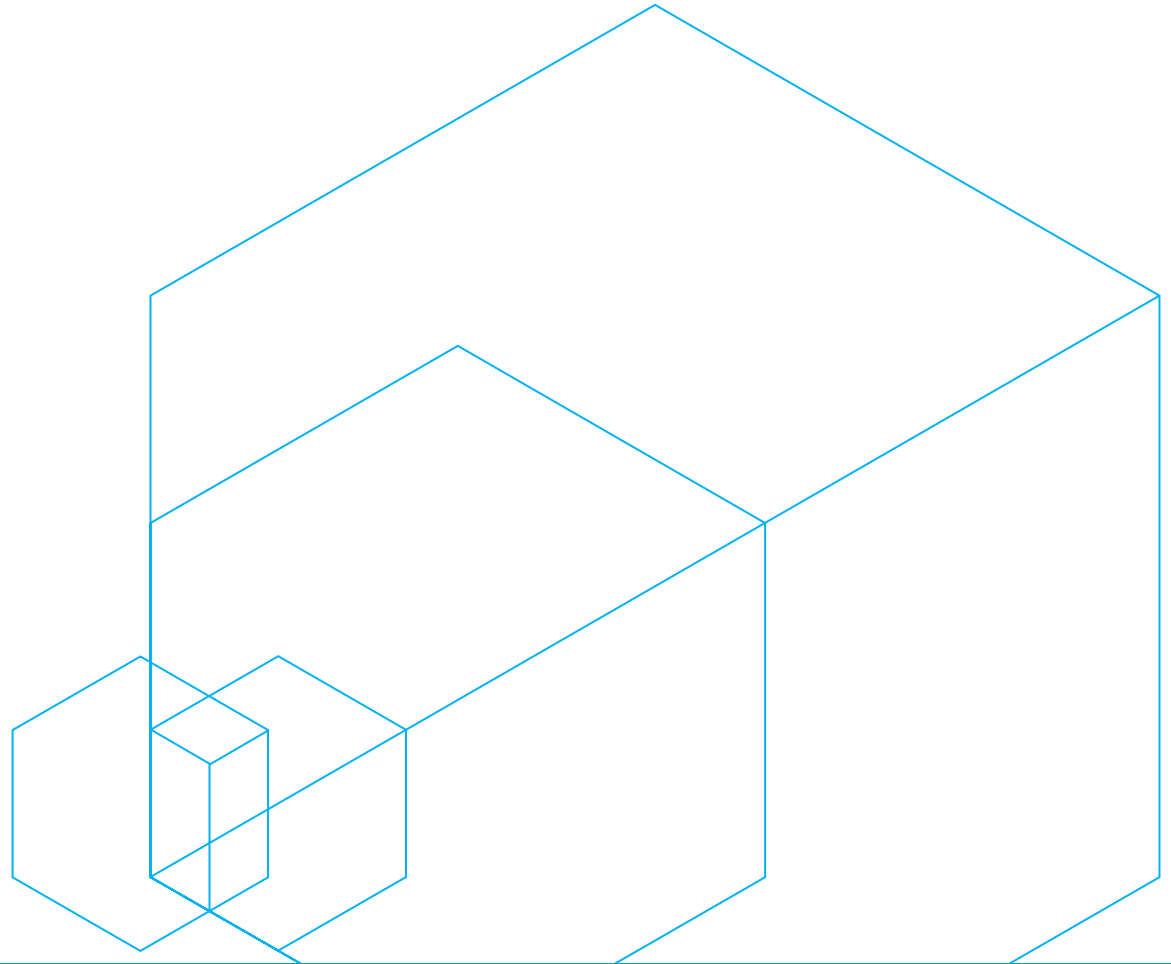
Software Bill Of Materials

Software	Name
Operating System	RHEL 7.2
Middleware	IBM Spectrum Conductor with Spark (integrated version) IBM Spectrum Scale (replication = 3) IBM Spectrum Conductor with Spark IBM Spectrum Cluster Foundation
Spark	<i>Included in IBM Spectrum Conductor with Spark</i>
GPU (optional)	CUDA 7.5? (Note: Not automatically deployed by IBM Spectrum Conductor with Spark)

Thank you.



ibm.com/systems



Please note:

- IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion.
- Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision.
- The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract.
- The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.
- Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.