Apache MESOS™

VS.

Krishna M Kumar, Lead Architect, Huawei@Bangalore

Mesos & Yarn

- *Both Allow you to share resources in cluster of machines.*
- *Mesos can manage all the resources in your data center but not application specific scheduling.*
- *YARN can safely manage Hadoop jobs, but is not designed for managing your entire data center.*
- *Two use cases – Mesos for non-Hadoop & Yarn for Hadoop.*

Mesos vs. Yarn

Mesos Architecture

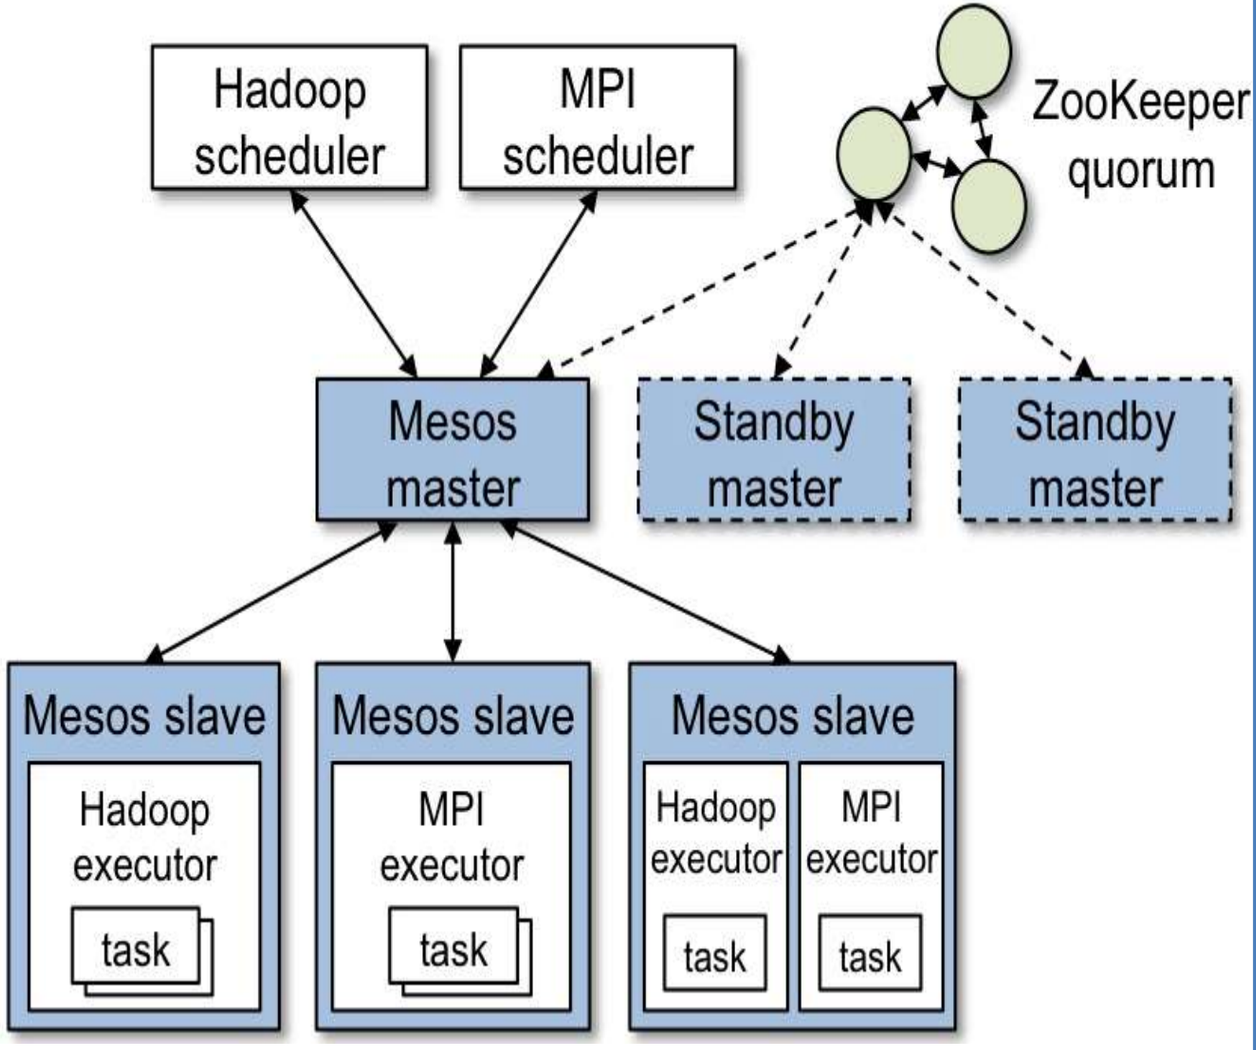

Here is an architectural view of YARN:

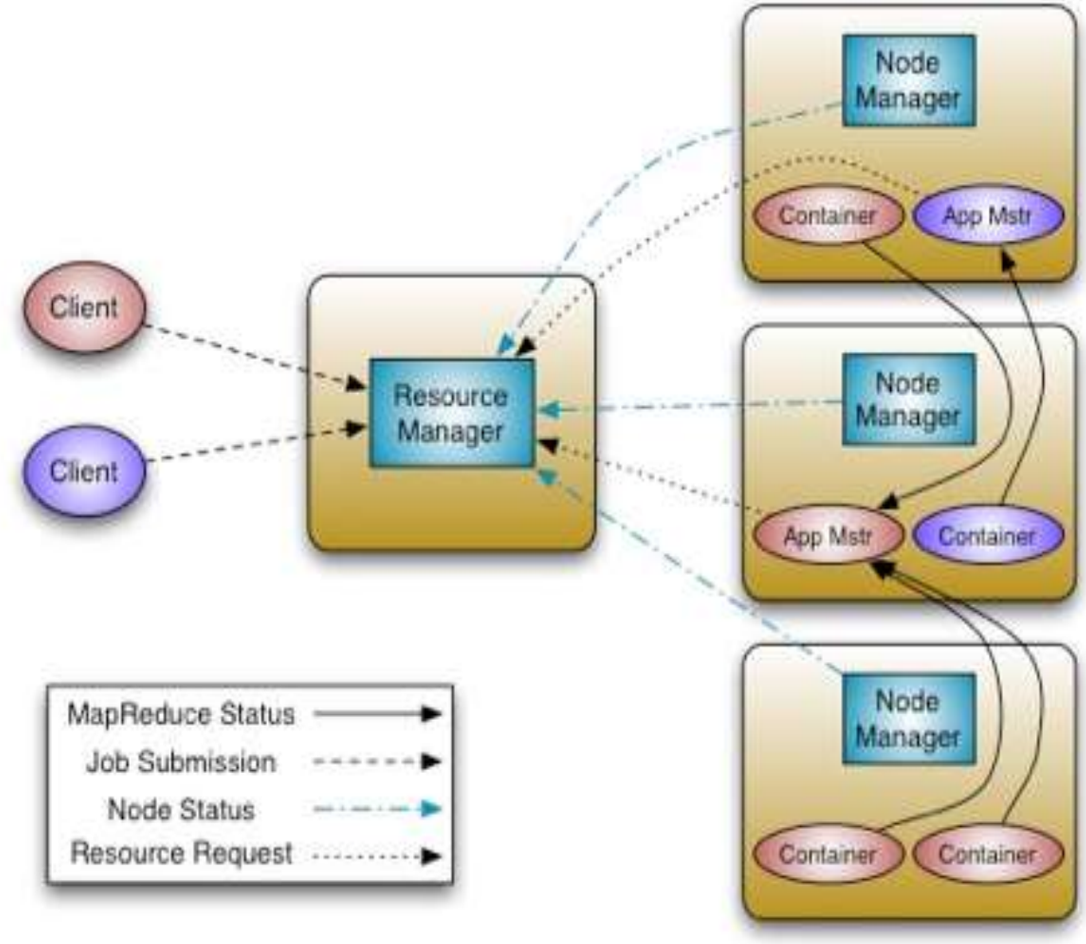

Mesos vs. Yarn

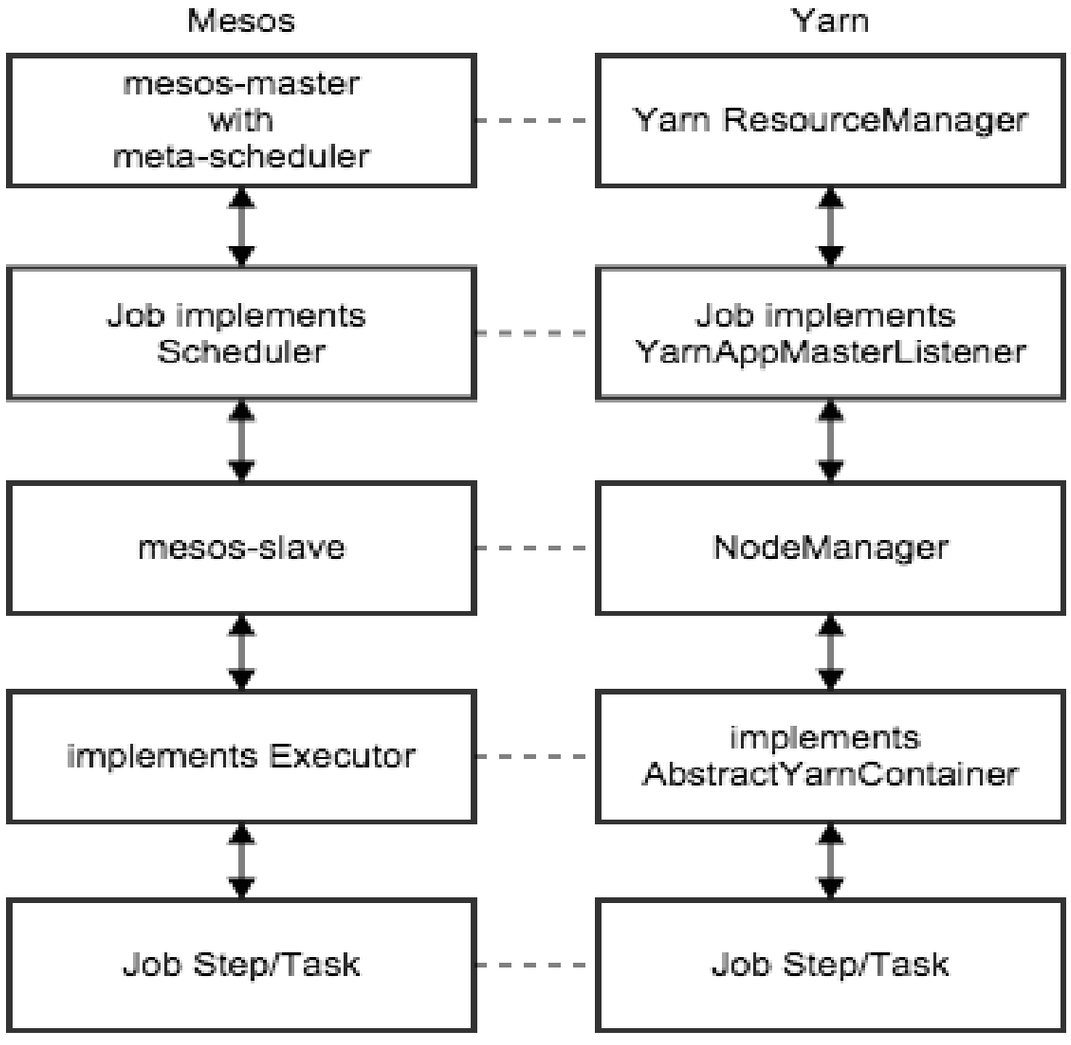

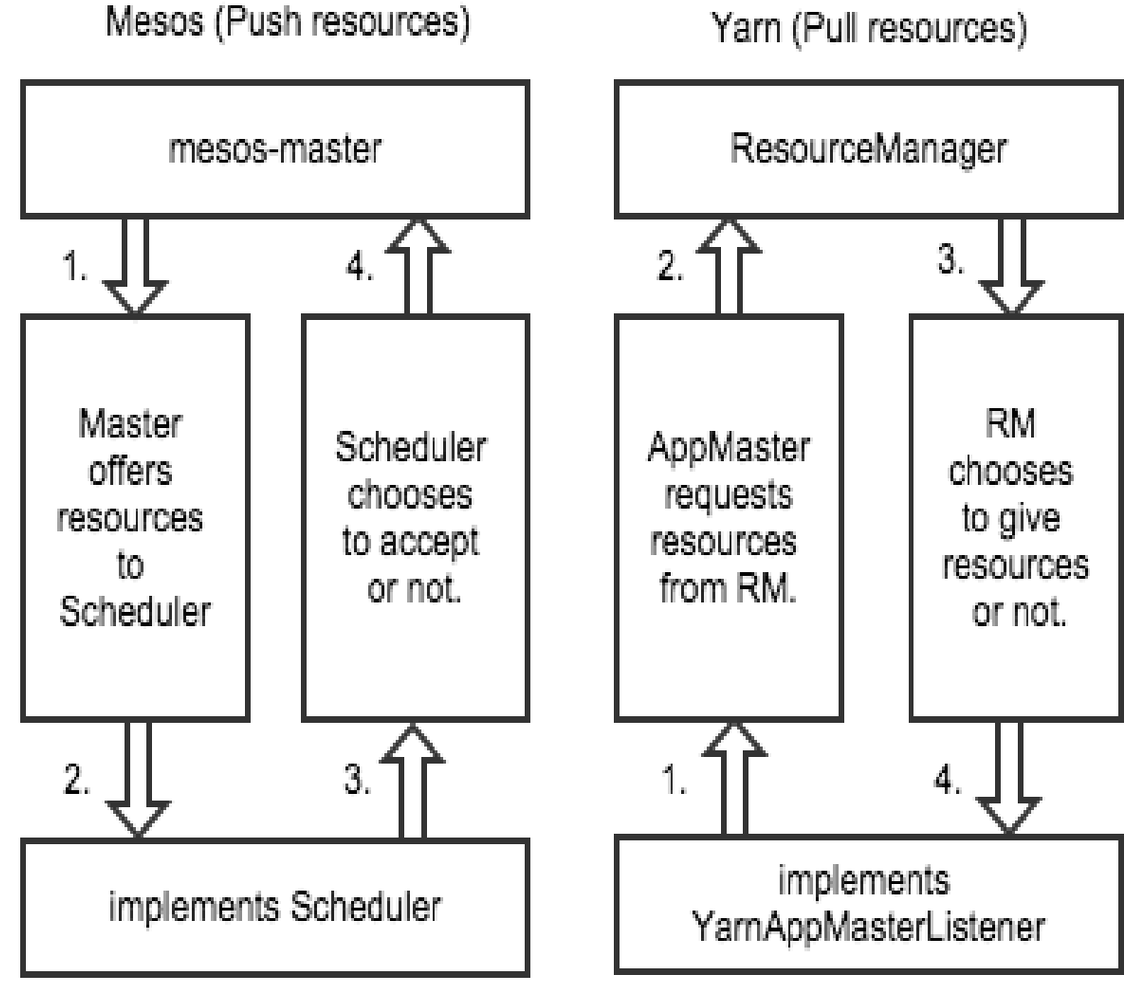

Myriad

Myriad – A open source software project is both a Mesos framework and a YARN scheduler that enables Mesos to manage YARN resource requests.

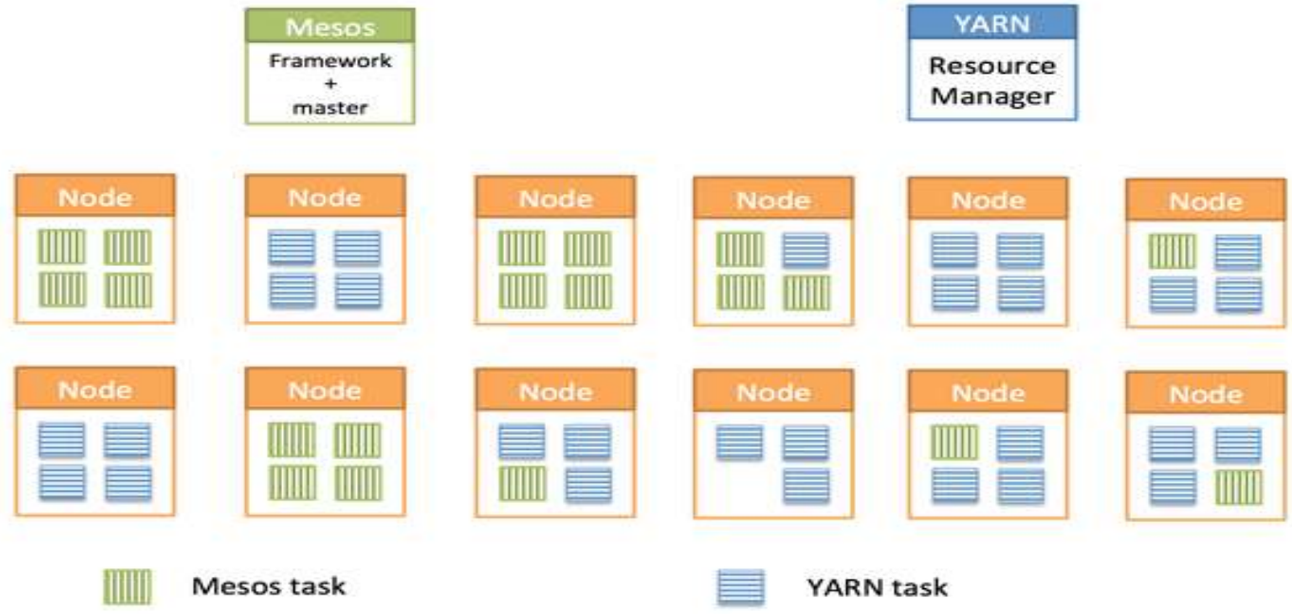

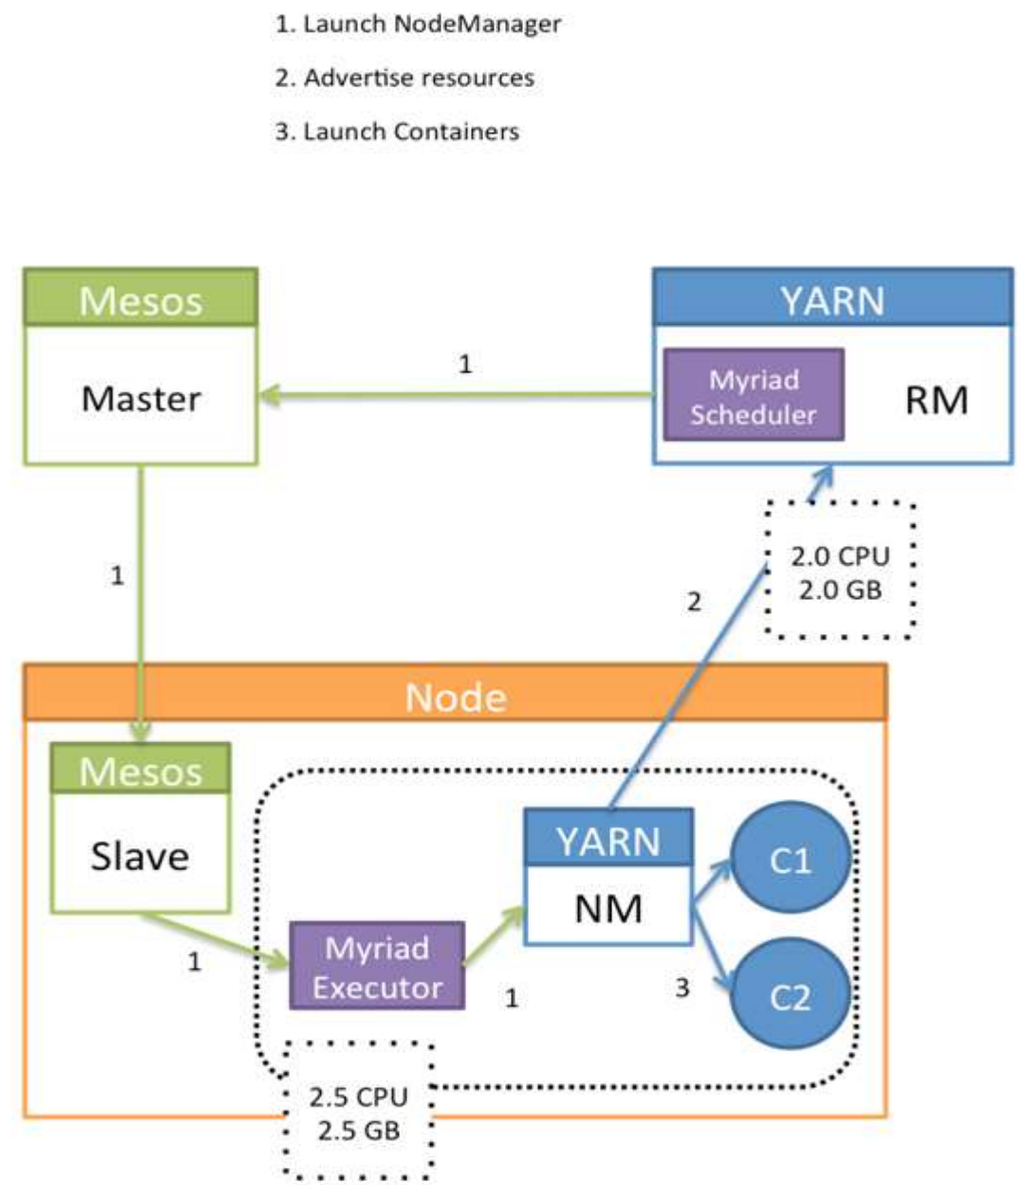

YARN on Mesos will bridge the world of Mesos and big data

Mesos	Yarn
Written in C++, good for time sensitive works	Written in Java, JVM based app
Memory & CPU scheduling – **Push based**	Mainly memory scheduling – **Pull based**
Use Linux Container groups	Use Simple Unix processes
Framework get Resource offer to choose – very minimal information as just needed	Framework ask a container with specification + preferences(like local). Lots of information passed
Core Mesos is lighter but one need to write a Framework	It's a Framework of its own and so its 3x code vs.Mesos
Mesos, you need to deal with the security	Yarn inherit Hadoop security
Mesos is general purpose scheduler for Data Center. Application writer deploy applications the way wanted	**Mainly exists on Hadoop world – it's a application scheduler.** Framework setup unix process/application
The Framework takes care the application specific items	Can supports clustered applications
Fault tolerance, app portability, etc. - the Framework has to deal with	Can enforce global constraints and local – so application can deploy to right place
High performance Actor Style Messaging passing	Hadoop RPC Architecture – direction piggy backs heart beat
Lower level abstraction	Can run Yarn on Mesos (Myriad)

Mesos	Yarn
Production Hardened at massive scale at Twitter & AirBnb	Designed to do MapReduce Scale, used for Hadoop scaling at Yahoo & Hortonworks
It's a two level scheduler- schedulers on multi tenant on the same cluster – better scaling capacity is there	Evaluate all the request and place the job. Resource Manager does the job. It's a monolithic scheduler
Mesos supported non-Hadoop things such as Google Kubernetes, Docker, Redis, Spark, etc. Supports Resource revocation, Pre-emption, dynamic reservation, optimistic offers	Good for batch jobs for long run times (e.g: Hadoop) & stateless batch jobs which can restart easily. Not designed for long running services, short lived queries, Web servers, SOA apps, real time workloads, state full services like DB or distributed File system.
Works well with legacy applications which can be scheduled in VM, Containers & BareMetal. Multiple Frameworks can be called as needed.	Legacy applications must be scheduled with customization.
Ben Hindman and the Berkeley AMPlab team worked closely with the team at Google designing Omega so that they both could learn from the lessons of Google's Borg and build a better non-monolithic scheduler. Written in UC Berkeley & project now with Mesosphere	Written by Yahoo/HortonWorks. YARN was created out of the necessity to scale Hadoop. Prior to YARN, resource management was embedded in Hadoop MapReduce V1, and it had to be removed in order to help MapReduce scale.

References

1) https://www.quora.com/How-does-YARN-compare-to-Mesos - with Arun Murthy's(Hadoop Founder) comments
2) https://www.oreilly.com/ideas/a-tale-of-two-clusters-mesos-and-yarn
3) https://www.mapr.com/blog/apache-mesos-vs-hadoop-yarn-whiteboard-walkthrough
4) https://mesosphere.com/blog/2015/02/11/yarn-on-mesos-big-data/
5) https://www.youtube.com/watch?v=aXJxyEnkHd4
6) https://engineering.twitter.com/university/videos/running-yarn-alongside-mesos
7) http://www.slideshare.net/payberah/mesos-43904525?qid=eb2d0eb2-c866-4826-819b-3cb405183231
8) http://www.slideshare.net/sameertiwari33/scheduling-on-large-clusters
9) http://blog.typeobject.com/a-quick-comparison-of-mesos-and-yarn
10) https://dzone.com/articles/workload-and-resource-management-yarn

Of course, some of these information was pulled directly from web without any modifications.