



Bursting Analytics and BI Workloads to a Public Cloud Provider in SAS Grid Manager

Gregg Rohaly
Global Business Development Manager - SAS Account
IBM Systems - Software Defined Infrastructure



Why? What is motivating Cloud Use...

SAS needs performance!

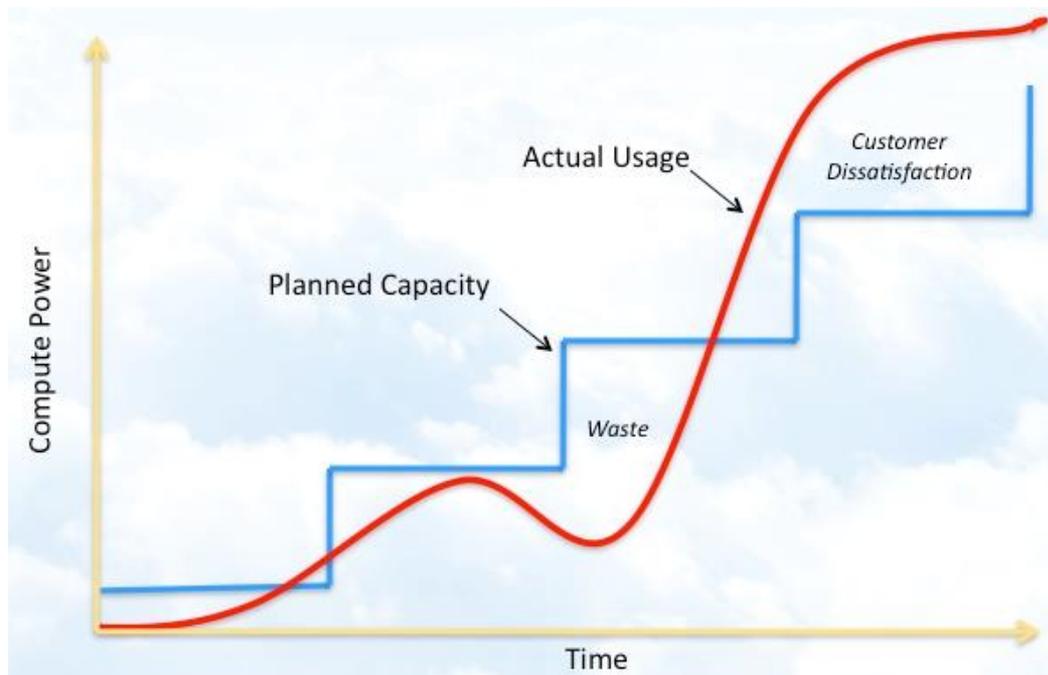
Major Trend HPC in the Cloud



- **64% of HPC sites run some jobs in public clouds**
 - Up from 13% in 2011
- **But only 7-8% of all their jobs (average)**
 - Not much changed since 2011
 - Public clouds are cost-effective for some jobs, but up to 10x more expensive for others
 - Key concerns: security, data loss
- **Private and hybrid cloud use is growing faster**
- **Big public clouds are going heterogeneous**
 - AWS with Ryft FPGAs, Google with NVIDIA GPGPUs

* From early 2017

Cap Ex versus Op Ex



What Problem We Are Trying To Solve

The Business' Need: Shorter term projects come up, the need for additional resources arises monthly/quarterly, and/or the cost of maintaining dev/test continues to put pressure on the purely on-premise model.

The Business' Pain Point(s): The core issues are buying and provisioning extra infrastructure quickly [easier said than done] or keeping additional resource at the ready, just in case.

The IT Solution: The hybrid cloud model permits a mix of on-premise resources with paying for the short- or long-term use of additional IT infrastructure and the associated compute capacity in the Cloud.

***Assumptions:** It is cost prohibitive to move your whole analytics environment to the Cloud. Certain data needs to be locally or country controlled for compliance/regulatory purposes. You do not TYPICALLY have pools of unused resources readily available.

The Answer! Now SAS Grid Manager can extend your usage model through bursting workload to additional capacity in the Public Cloud.

Note: Cloud provider performance may vary significantly. We will discuss later in the talk.

SAS Grid Manager Key Benefits

Capability

Why it Matters

Centralized/Shared Workload Management



Effectively manage jobs and users, prioritization, compliance/auditing

High Availability



Avoid user or service disruption, rolling maintenance

Distributed Processing



Improved performance, meet changing demands, grow incrementally

Leverage Commodity Hardware



Reduce Costs

Distributed Analytics Computational Platform



Grid-Processing
(SAS Grid Manager)



In-Memory Processing
(SAS Viya CAS Engine)



Workload
Management

Dynamic: Prioritization | Pre-emption | Thresholds
Manages workloads @ Job Level

Relies on OS for concurrent activity management
Works as one memory block of entire cluster or subset nodes

Integration Point: SAS Grid manages and throttles concurrent SAS jobs that make use of CAS engine actions

High
Availability

HA for Infrastructure: Multiple machines | LSF daemons
HA for Batch Jobs: Re-queue & checkpoint/restart
HA for Critical Services: use the Platform LSF Enterprise
Grid Orchestrator (EGO)

HA for Infrastructure: Back-up Controller and Failure-resistant Node architecture
HA for data and procs: execution w/ data duplication
HA Sessions: Separate CAS user session | isolation

Parallelization

At Job/Program level
Users can submit multiple concurrent || jobs

At Procedure level (what is CAS enabled for scale)-
based on succession of input/output/formats

Integration Point: The individual || Jobs on Grid can contain CAS enabled procedures that further scale in MPP

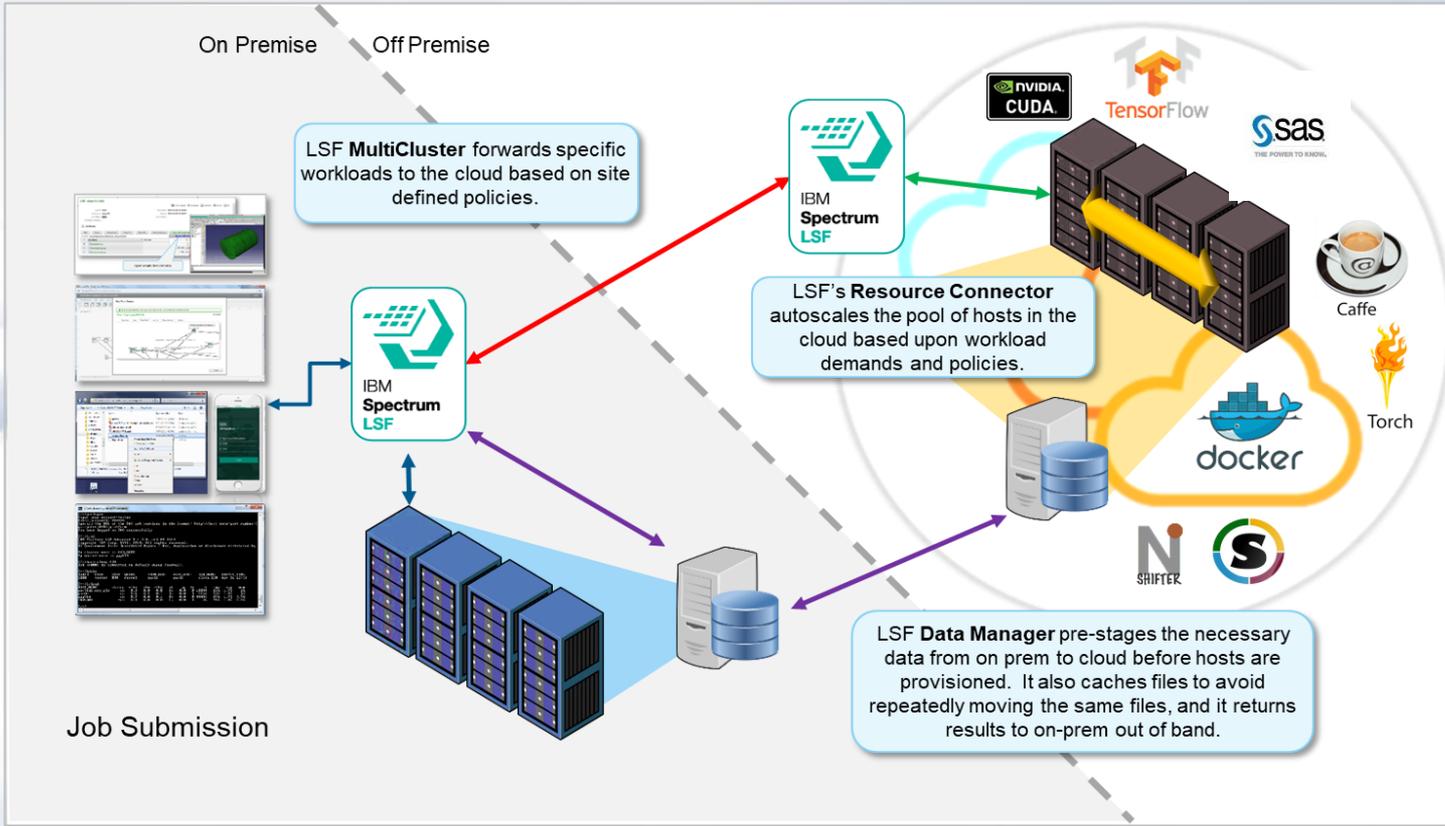
Best Use

Lots of users, lots of SAS applications
|| Data Manipulation and Model Training Jobs

Large Scale Data Manipulation and Model Training
Interactive | Scalable Visualization

Integration Point: Scaling and Bursting of || Model Training Jobs

IBM Spectrum Computing's HPA/C Hybrid Cloud Capability



Amazon



IBM Cloud



Microsoft



Google

IBM Spectrum LSF Data Manager [DM] Functionalities

Data Management

- DM is cache aware. If multiple jobs share same data source, DM only transfers the data once.
- DM manages the data life cycle; it deletes the cached data without reference automatically.

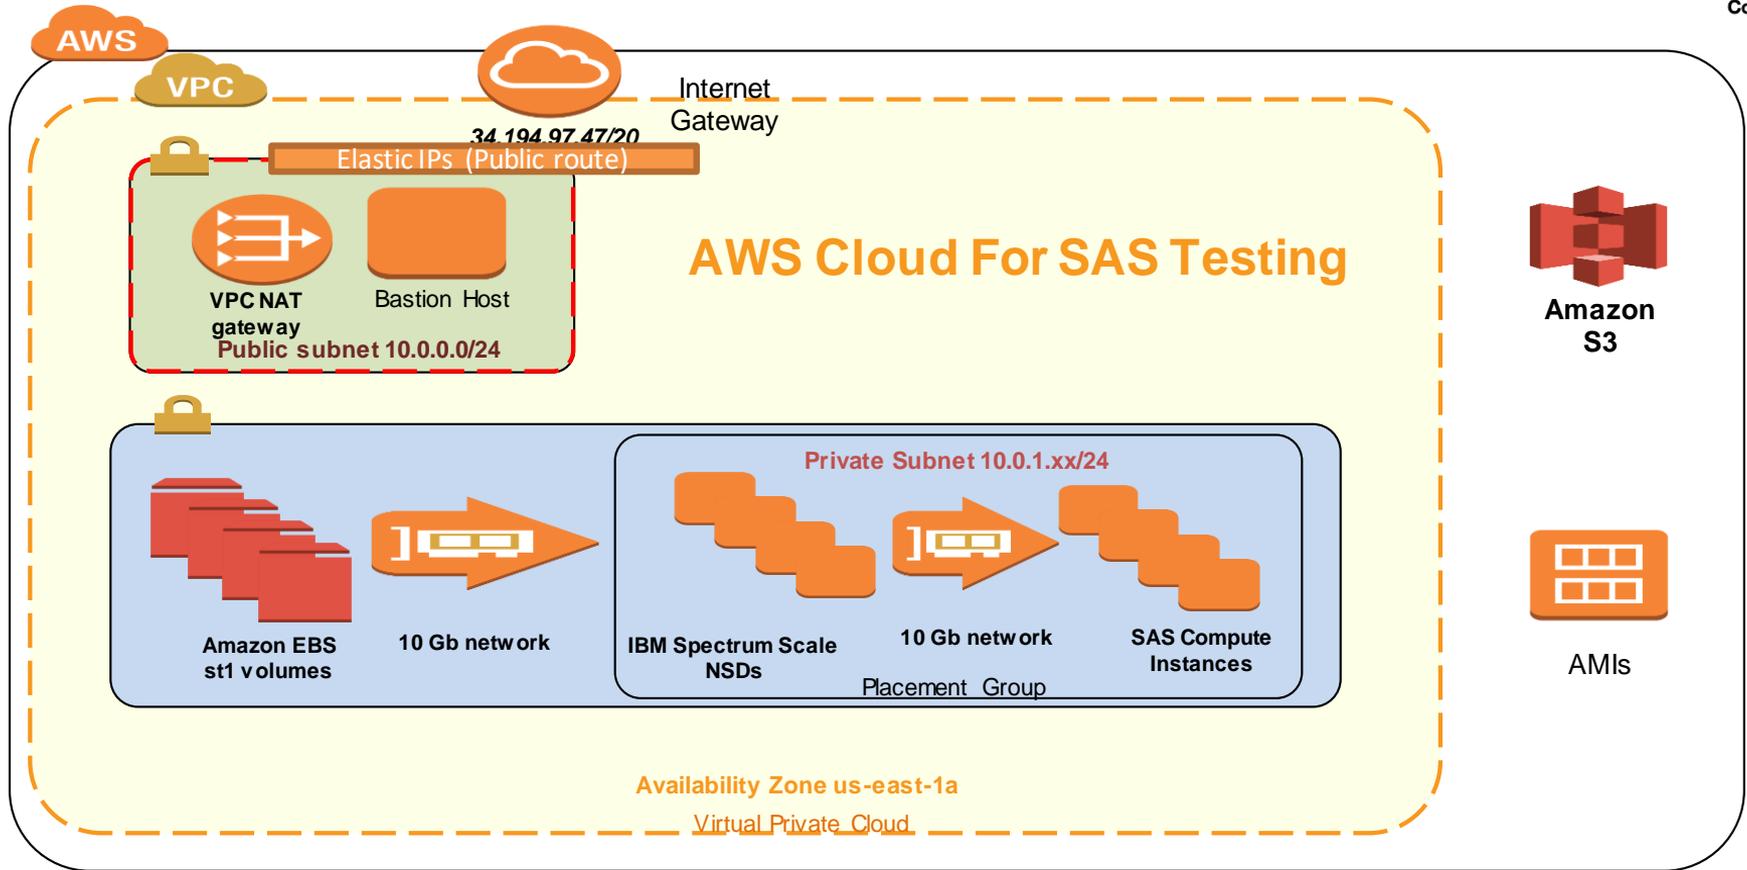
Smart Data Staging

- DM decouple data transferring from job's allocation. It transfers data when job is queuing, instead of when job is starting. This avoid wasting resources.
- Provides tools "bstage" to stage data between "staging area" (in general, it's a shared storage) and local disk

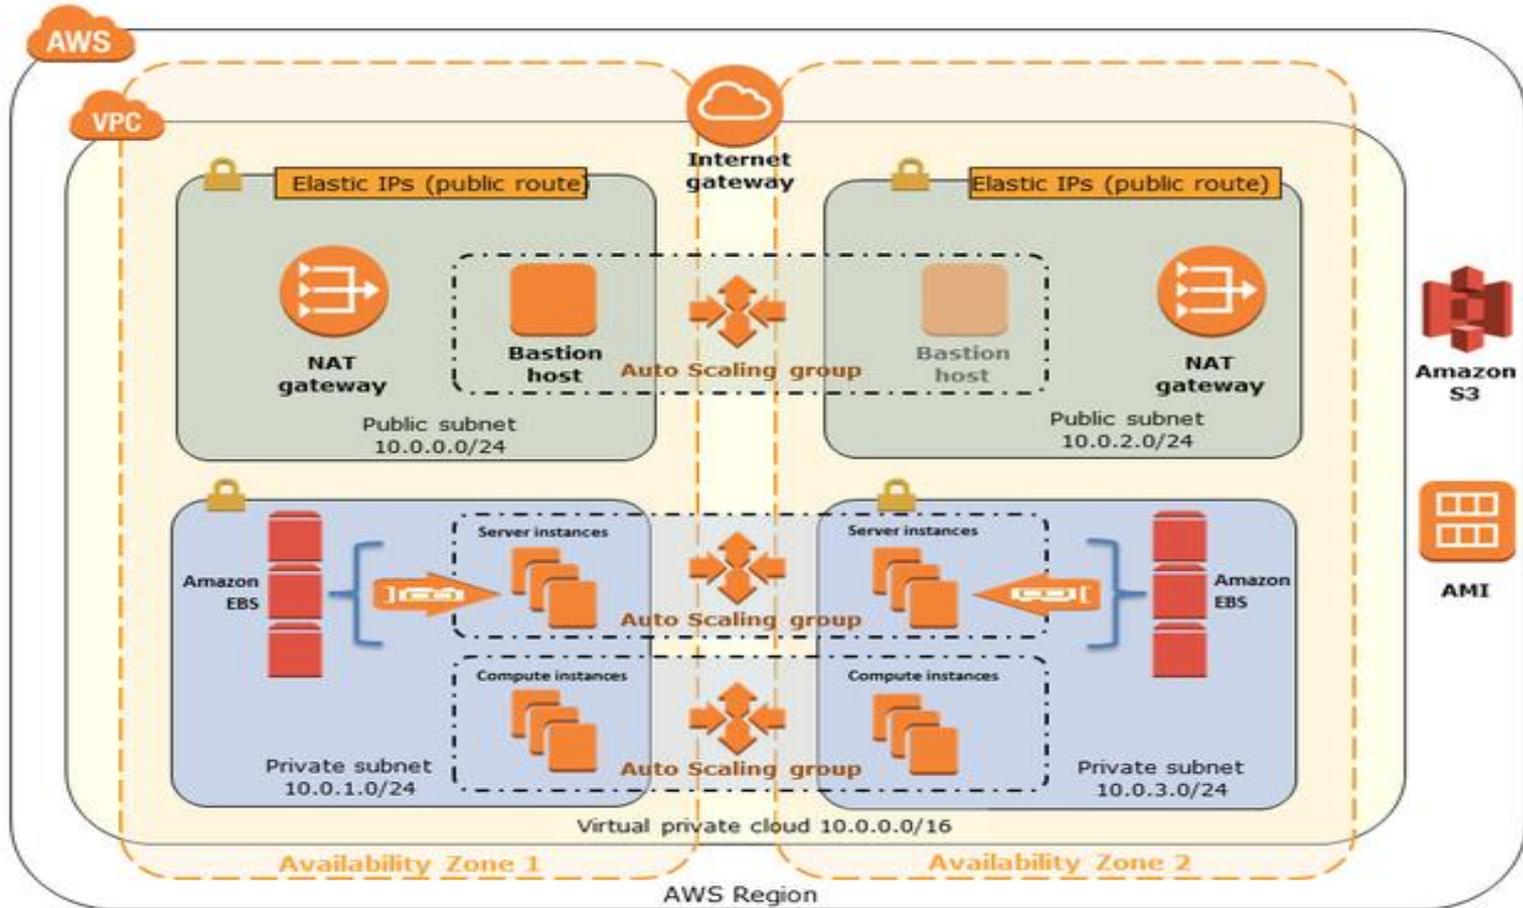
Data Transfer

- DM supports customizing data transfer tools, like scp, rsync, ascp etc.
- DM manages the data transferring process; it controls the maximum concurrent data transfers.

SAS Grid Manager on AWS Architecture [Bursting to Location]



SAS Grid Manager on AWS Architecture [Bursting to Location w/ DR]



Performance Heavily Depends Upon the Cloud Providers Plumbing

- The governance of the 10Gbit network fabric in AWS, and other Cloud providers, will constrain the overall IO potential. This is due to the NIC to EBS limit of 875 MB/s throughput of the best performing systems today in 10Gbit fabric. This equates to a potential IO of 54.6 MB/s/core of the 16 core systems tested.
- SAS, as of this writing, recommends at least 125 MB/s/core of throughput or more than 2.28 times the current AWS, and other typical Cloud provider, limits. Understanding this point is key and essential!
- It is important to note that, when the VPC was created, the instances were defined using an AWS placement group in one availability zone only. This is important for two performance reasons:
 - ✓ Placement groups insure all the nodes are in close proximity to each other for more predictable network throughput.
 - ✓ A placement group can span availability zones within a VPC and AWS recommends doing so for High Availability (HA). Spanning availability zones automatically limits the network performance from 10 Gbit/second to 5 Gbit/second.

SAS Grid Manager Bursting Control and Decision-making

SAS Grid Manager has two places to configure the bursting behavior: 1) the queue level and 2) at the LSF resource connector policy level.

First, for the queues which are cloud-enabled, the primary decision is made if the workload demand threshold is met for cloud bursting to occur. If true, then control passes to the LSF resource connector and the pre-defined policies that are in place. It will evaluate how many instances should be launched to join the cluster. Keep in mind that precise and restrictive policies can be defined to fine tune the bursting behavior. Some [but not all] of these examples include:

- Configuring a threshold on when to launch instances
- Throttling the rate of launching instances
- Maximum instances to request
- Setting policies to reclaim the instances

The Good News – Bursting Can Benefit the Business

- This is doable AND businesses are taking advantage of bursting to public cloud available resources NOW!
- Quick response times in obtaining additional resources should permit quicker results even with certain current Cloud provider performance limitations. WARNING: It's the Cloud, but not everything is well suited for it at this time.
- Security, networking, storage, business continuity, etc models exist in Cloud providers such that an enterprise can be resilient using this approach.
- SAS Grid Manager has this core capability available as of SAS 9.4M5. We are looking for Clients to be initial adopters.

Q&A and Discussion

