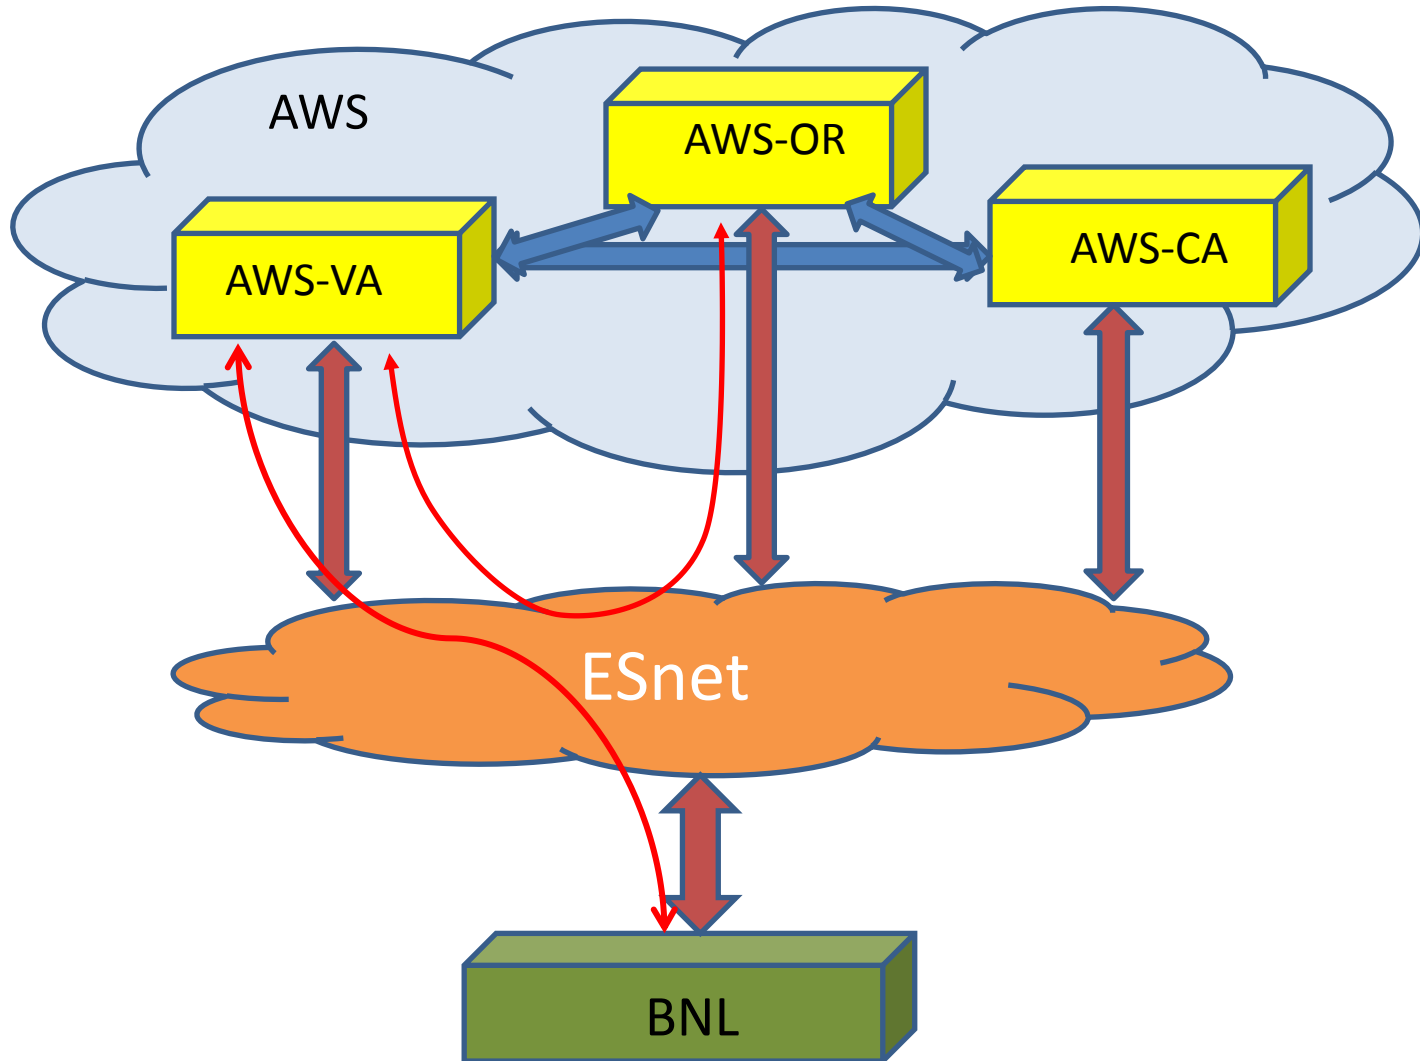


BNL/Amazon AWS Pilot

Michael Ernst, BNL

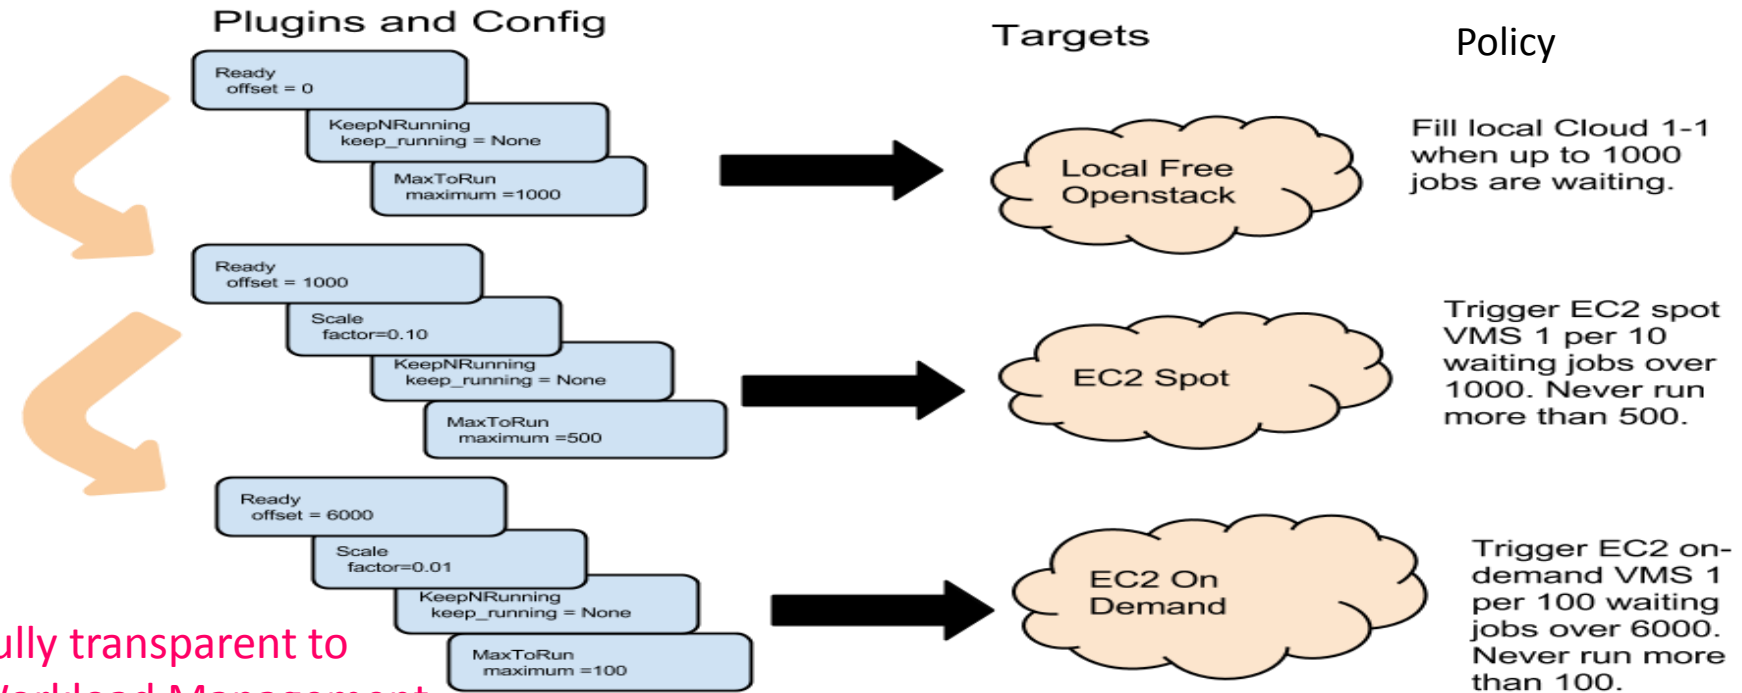
- Who?
 - AWS, T1/BNL and ESnet
- Objectives
 - To evaluate technical feasibility of using AWS resources at a scale of 50k-100k concurrent jobs to run ATLAS production and analysis workflows
 - CPU added “elastically” to dedicated resources
 - Added to existing HTCondor pool by APF-based Cloud Scheduler
 - Dynamic resource provisioning (EC2 Spot) at scale on short timescale
 - Petascale storage in the Cloud to host objects, datasets produced in the Cloud and replicated to the Cloud for production and analysis workflows
 - Create & operate an ATLAS compatible SE in the Cloud based on S3
 - » Needs modifications to the PanDA pilot – Done by Wen
 - » Needs S3 support in FTS3 – Will be done by FTS3 developers
 - AWS Network integration with ESnet
 - Goal: Wave AWS Egress fee and avoid AWS inter-region network fee
 - Make AWS facilities a directly connected ESnet site (replace peering)
 - Use network virtualization to make AWS resources part of BNL/ATLAS network (use Amazon SDN capabilities)

Networking




Using Cloud Resources effectively: A Policy-based Cloud Scheduler, developed at RACF

Example: Cascading Cloud Targets

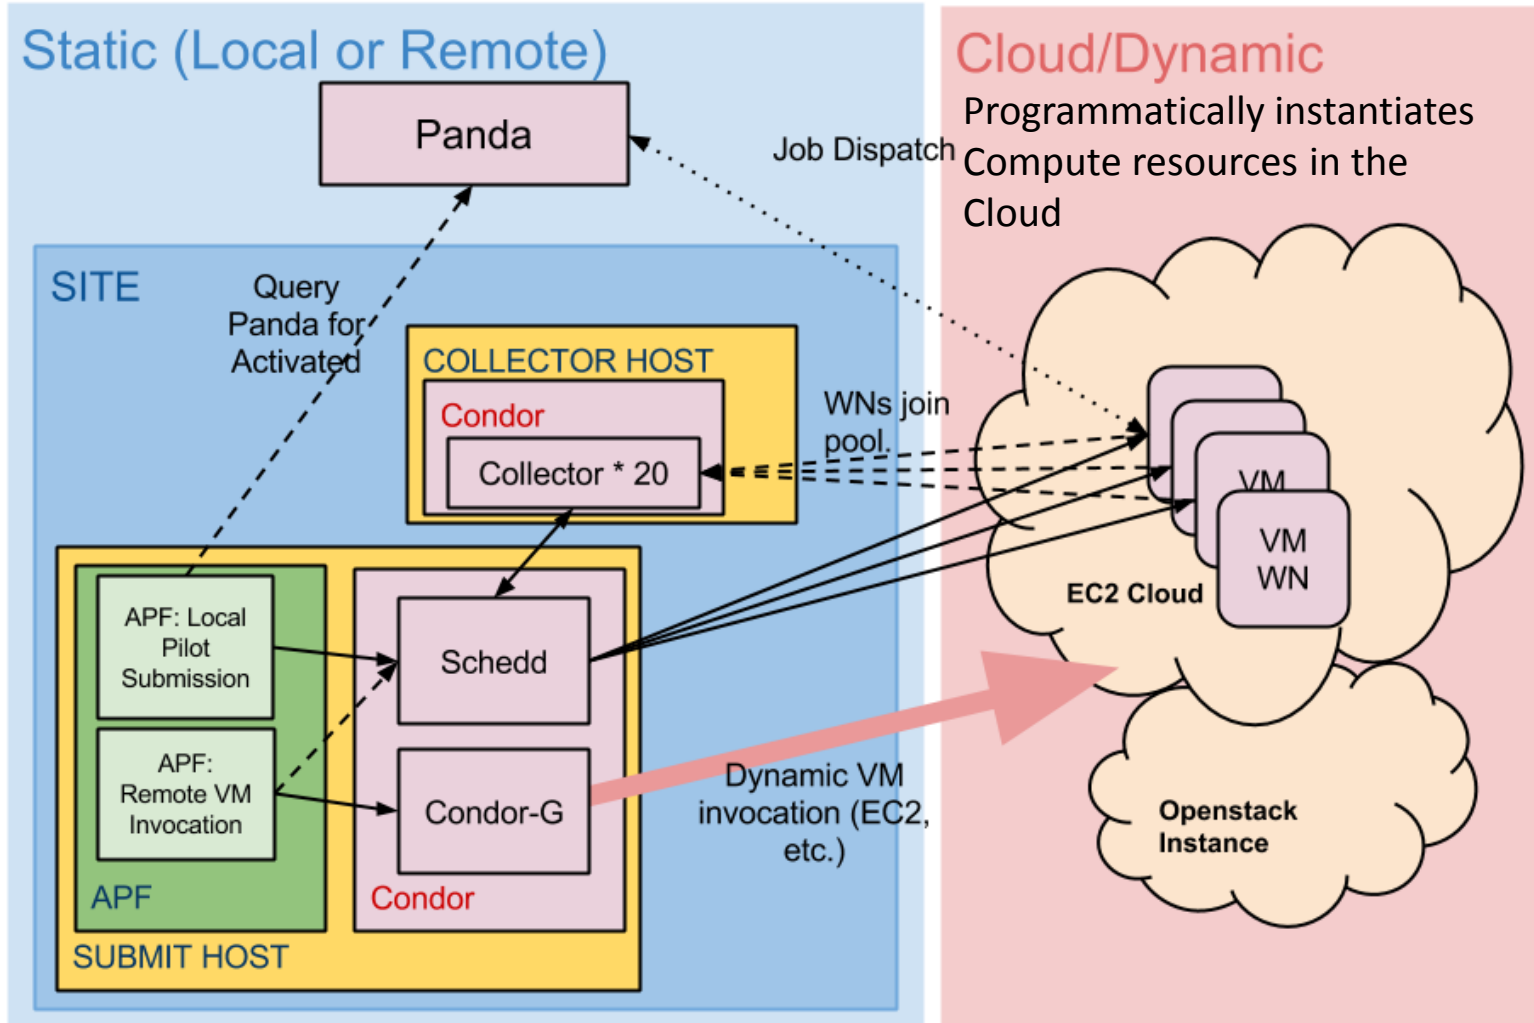


Fully transparent to Workload Management System (e.g. PanDA), Elastically expands Pool of Compute Resources according to user-defined policy




Demand-driven, policy-based programmatic instantiation and contraction of cloud resources

Elastic Cluster



John Hover, BNL



Designed to serve

- Peak demands
- New users w/o dedicated res.


BNL/Amazon AWS Pilot

- Objectives (cont'd)
 - Monitor and analyze the usage pattern of all resources
 - CPU, Storage and Network
 - Understanding the usage pattern
 - US ATLAS Facilities: Understand to what extent AWS resources are usable to serve peak demand
 - AWS: Understand in detail how their resources are used by a science community and, based on this, will be in the position to develop a business model that will (hopefully) be attractive to us (and other communities)
 - Amazon has recently formed a Scientific Computing Group
 - » Members of that group are an integral part of the pilot project. They get and manage the AWS resources (human and facilities) we need to carry out the pilot
 - » AWS provides grant for resources needed according to a pre-defined project plan (quantities and workflows)
 - ESnet Director and his engineering team are very interested in the Pilot as it will provide guidance as to how the research community may use commercial Cloud resources

	Name	Begin date	End date
Done	AWS account setup	8/1/14	8/7/14
	Pilot project management (user side)	8/1/14	11/20/14
	Setup Job Submission Infrastructure	8/4/14	8/18/14
	Prepare Storage Element at BNL	8/4/14	8/10/14
In progress	Setup perSONAR instances	8/11/14	8/19/14
	Characterize network links AWS <-> BNL	8/15/14	8/19/14
	Prepare ATLAS Workload	8/4/14	8/18/14
	Verify completion of above tasks	8/19/14	8/20/14
	MS 1: Prep work for use case 1 completed	8/20/14	8/20/14
	Run sample jobs for each workflow	8/21/14	8/24/14
	Run Simulation at scale	8/25/14	8/31/14
	Run Sim + Pile at scale	9/1/14	9/3/14
	Run Sim + Pile + "Derivation" at scale	9/4/14	9/7/14
	Run Sim + Pile + Derivation + RAW Reco at scale	9/8/14	9/14/14
Use ES for Sim Prod	Collect processing, storage & network stats	8/21/14	9/14/14
	MS 2: Exec of ATLAS workflows complete	9/15/14	9/15/14
	Use case 2: WLCG Storage Element in AWS	9/15/14	9/28/14
	Configure ATLAS DDM Endpoint	9/29/14	9/29/14
	Make Subscriptions against AWS Endpoint	9/30/14	10/7/14
	Run a few sets of Analysis Jobs against data o...	10/8/14	10/15/14
	Collect processing, storage & network stats	9/30/14	10/15/14
	MS 3: Eval of AWS Storage Element complete	10/16/14	10/16/14
	Set up AWS disk cache w/ EBS	10/17/14	10/26/14
	Run a few sets of Analysis Jobs against AWS R...	10/27/14	11/1/14
Pilot Task List	Analyze Accounting Data	11/3/14	11/18/14
	MS 4: End of pilot - all results documented	11/19/14	11/19/14

BNL -> AWS File Transfer Performance

S3cmd tool for Amazon Simple Storage Service
(S3) multi part size 100MB



10G connected Proxy Server reading from BNL/T1 prod SE via NFS4,
writing to AWS w/ S3 protocol