High Throughput Computing for Astronomers

Mats Rynge
Information Sciences Institute, University of Southern California

rynge@isi.edu

USC Viterbi
School of Engineering
Information Sciences Institute

Outline

- High Throughput Computing
  - Condor
    - DAGMan
      - Pegasus Workflow Management System
        - Periodogram Workflow / Open Science Grid
          - Galactic Plane Workflow / XSEDE
            - Cloud

Why High Throughput Computing?

For many experimental scientists, scientific progress and quality of research are strongly linked to computing throughput. In other words, they are less concerned about instantaneous computing power. Instead, what matters to them is the amount of computing they can harness over a month or a year --- they measure computing power in units of scenarios per day, wind patterns per week, instructions sets per month, or crystal configurations per year.

Slide credit: Miron Livny

High Throughput Computing is a 24-7-365 activity

FLOPY ≠ *(60*60*24*7*52)*FLOPS*

Slide credit: Miron Livny

HTC Toolbox

Condor

Matchmaking and scheduler for HTC workloads

DAGMan

Directed Acyclic Graph Manager – workloads with structure

Pegasus

Workflow Management System

Condor

Users submit their serial or parallel jobs to Condor, Condor places them into a queue, chooses when and where to run the jobs based upon a policy, carefully monitors their progress, and ultimately informs the user upon completion.

Job queuing mechanism

Scheduling policy

Priority scheme

Resource monitoring

Resource management

DAGMan

Directed Acyclic Graph Manager

Handles tasks with dependencies

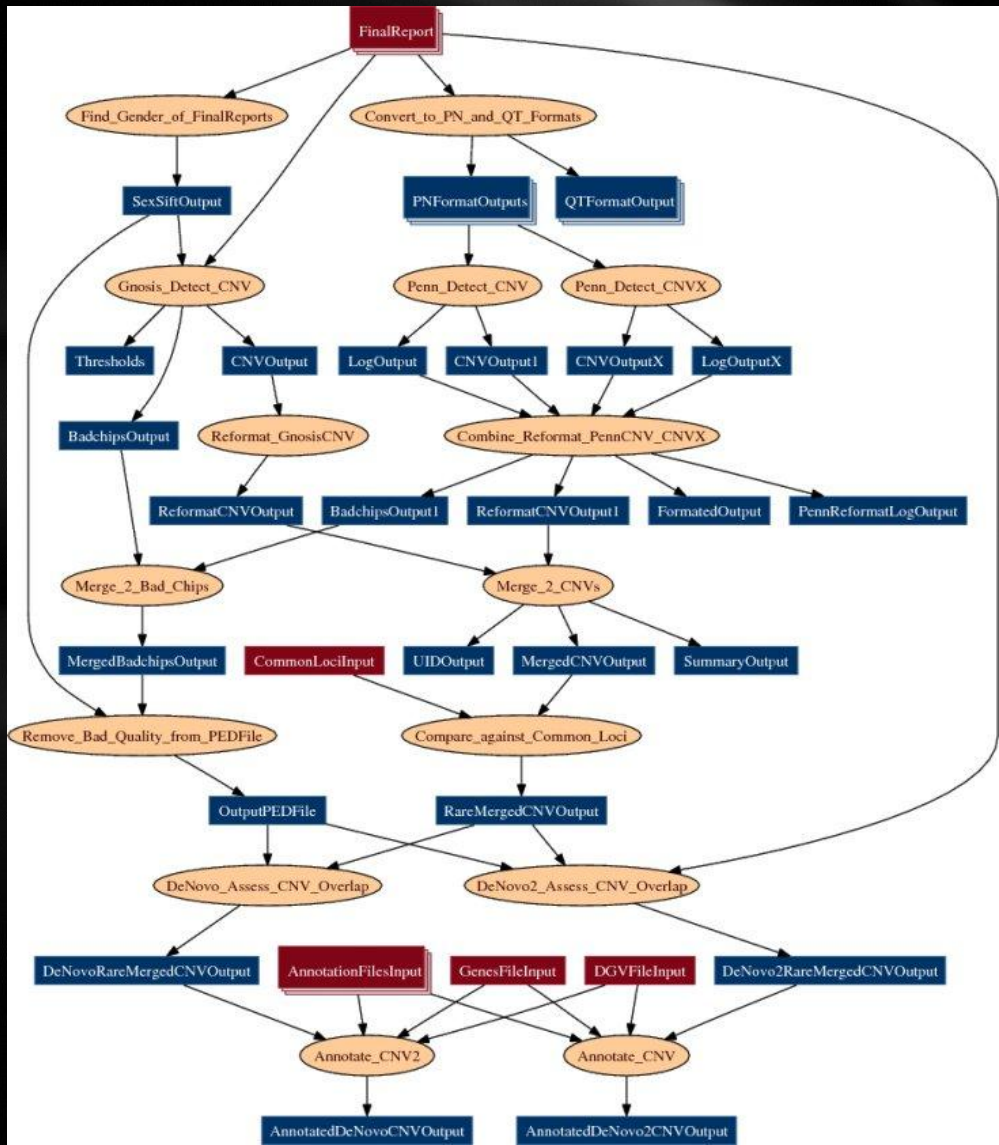

Pegasus Workflow Management System

Abstract Workflows - Pegasus input workflow description

- Workflow "high-level language"
- Only identifies the computation, devoid of resource descriptions, devoid of data locations

Pegasus

- Workflow "compiler" (plan/map)
- Target is DAGMan DAGs and Condor submit files
- Transforms the workflow for performance and reliability
- Automatically locates physical locations for both workflow components and data
- Provides runtime provenance

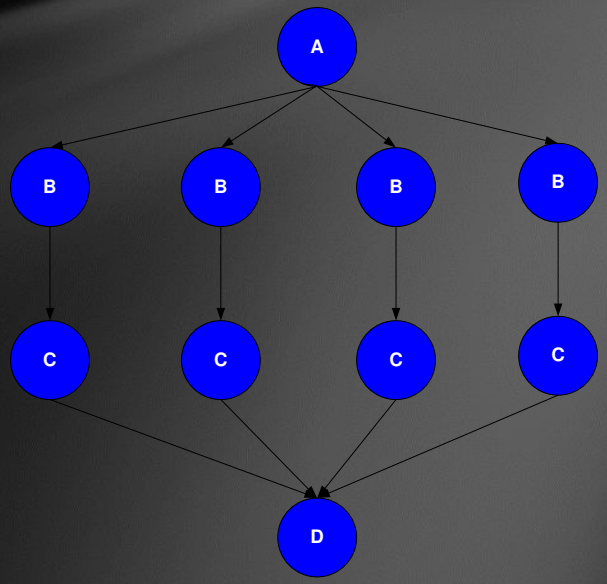

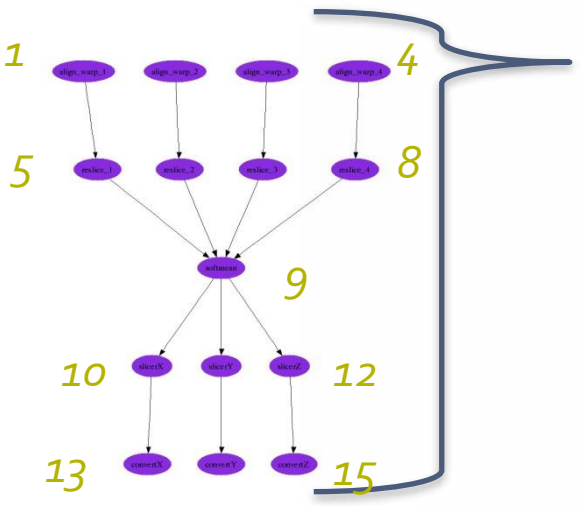

13 data stage-in nodes

11 compute nodes (4 reduced based on available intermediate data)

8 inter-site data transfers

14 data stage-out nodes to long-term storage

14 data registration nodes (data cataloging)

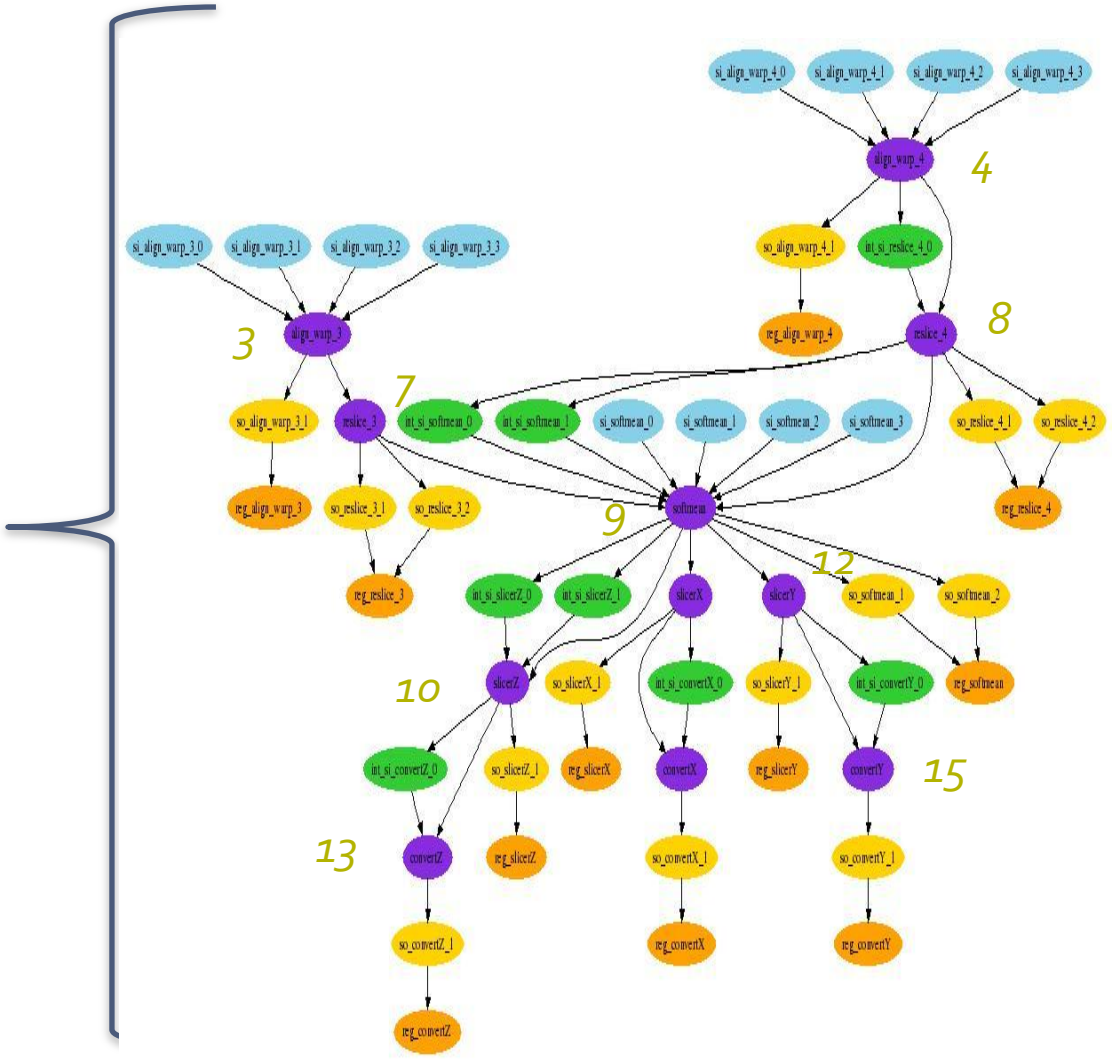

If You can draw….

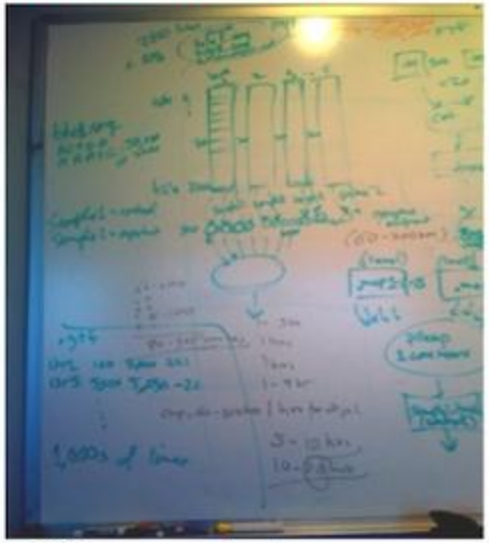

We can make it run….

pegasus

Epigenomics, Ben Berman, USC

PSOCI (Solar Fuels) Jeff Tillson, RENCI

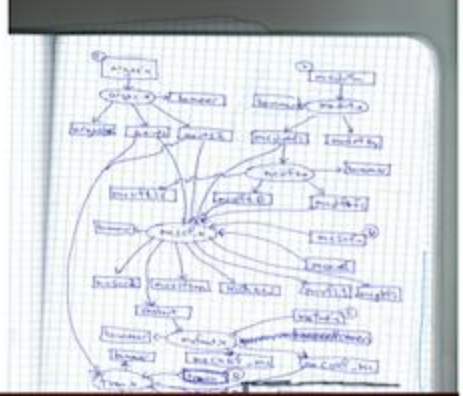

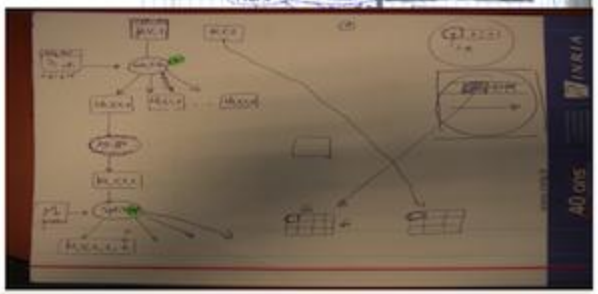

Helio-seismology, Laurent Gizon, Max Planck

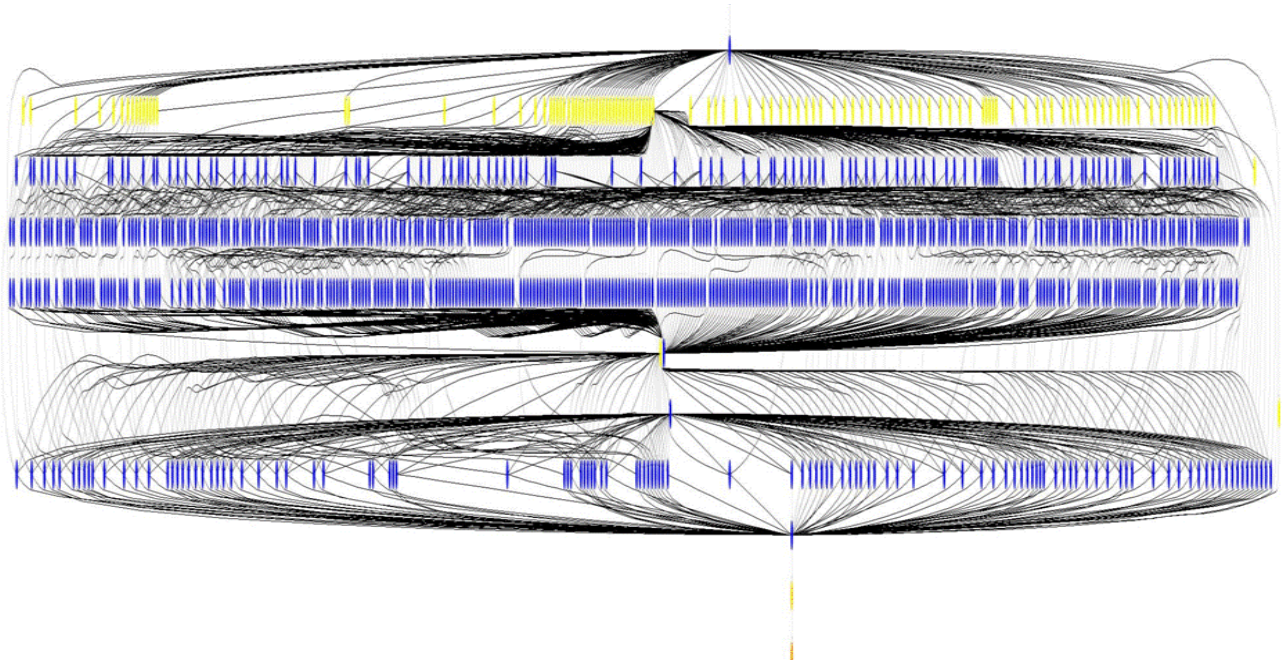

Periodogram Workflow

- 1.1 million total tasks
- 108 sub workflows
- Input: 323 GB
- Outputs: 2 TB
- 100,000 CPU hours

- Wall time based job clustering
  - Simple binning
  - Target: 1 hour

Scientific goal is to generate an atlas of periodicities of the public Kepler data. The atlas will be served through the NASA Star and Exoplanet Database (NStED) , along with a catalog of the highest-probability periodicities culled from the atlas.

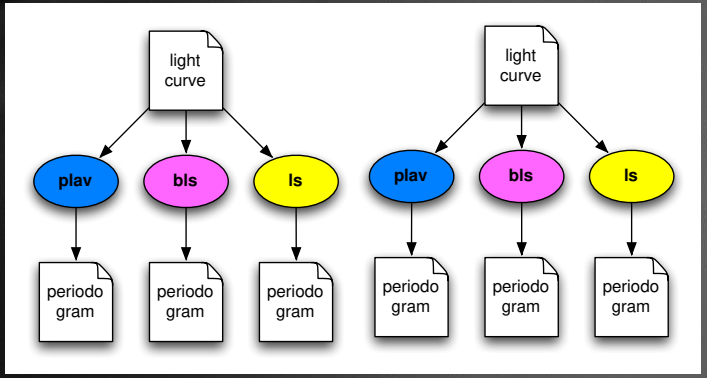

Periodogram Jobs Running on the Open Science Grid

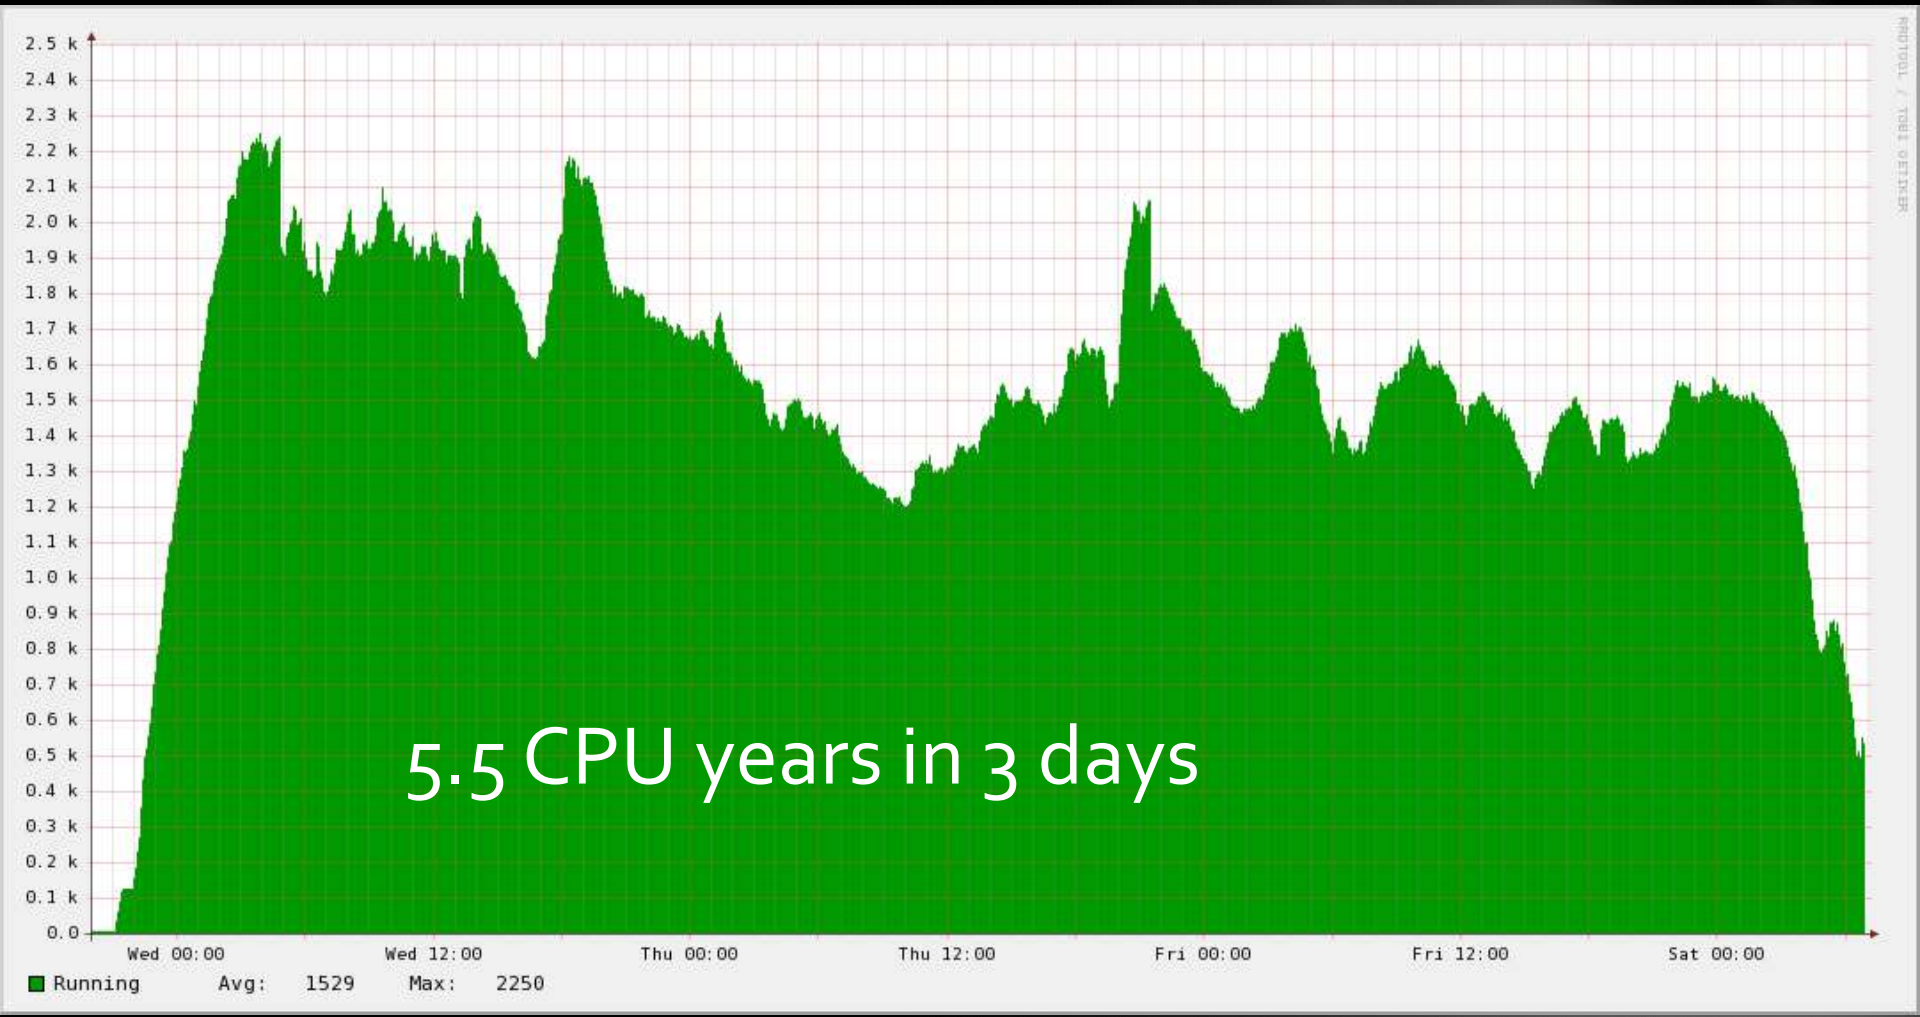

The Open Science Grid

A framework for large scale distributed resource sharing

addressing the technology, policy, and social requirements of sharing

OSG is a consortium of software, service and resource providers and researchers, from universities, national laboratories and computing centers across the U.S., who together build and operate the OSG project. The project is funded by the NSF and DOE, and provides staff for managing various aspects of the OSG.

Brings petascale computing and storage resources into a uniform grid computing environment

Integrates computing and storage resources from over 100 sites in the U.S. and beyond

The Evolution of OSG

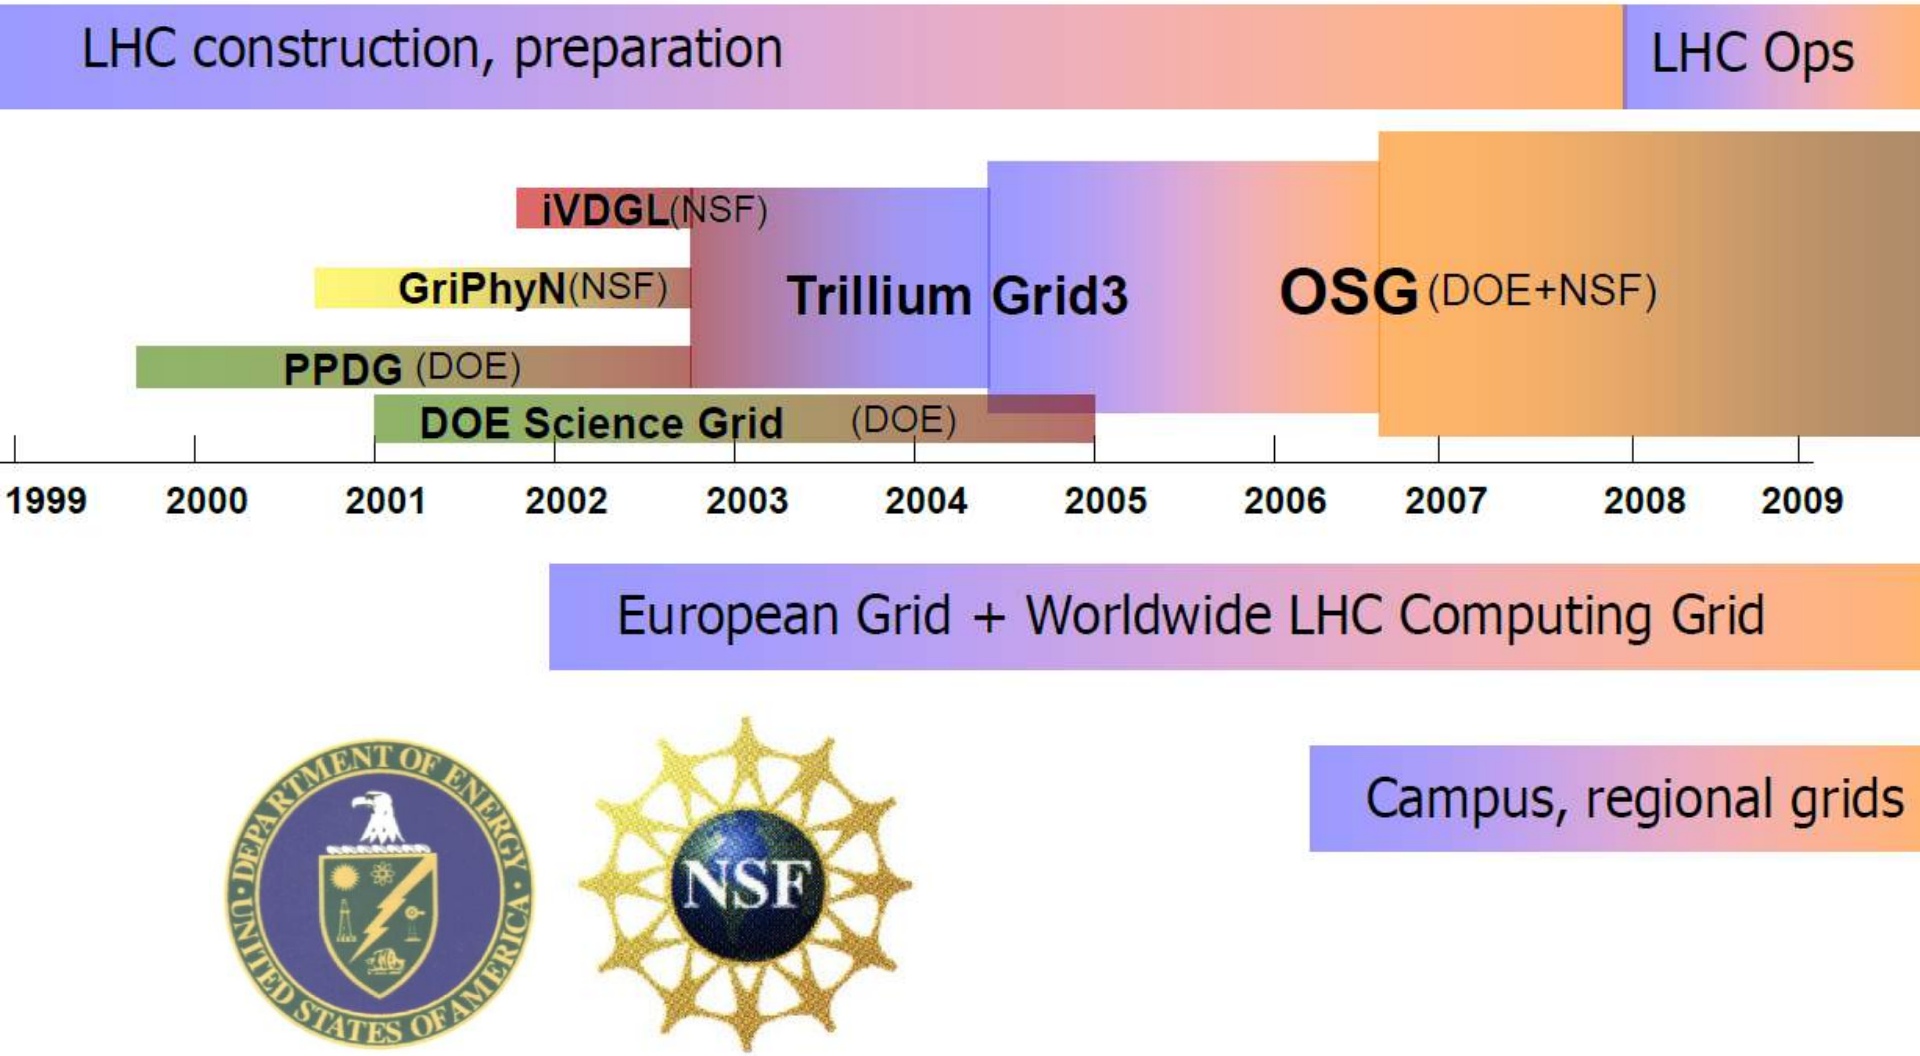

Using OSG Today

- Astrophysics
- Biochemistry
- Bioinformatics
- Earthquake Engineering
- Genetics
- Gravitational-wave physics
- Mathematics
- Nanotechnology
- Nuclear and particle physics
- Text mining
- And more…

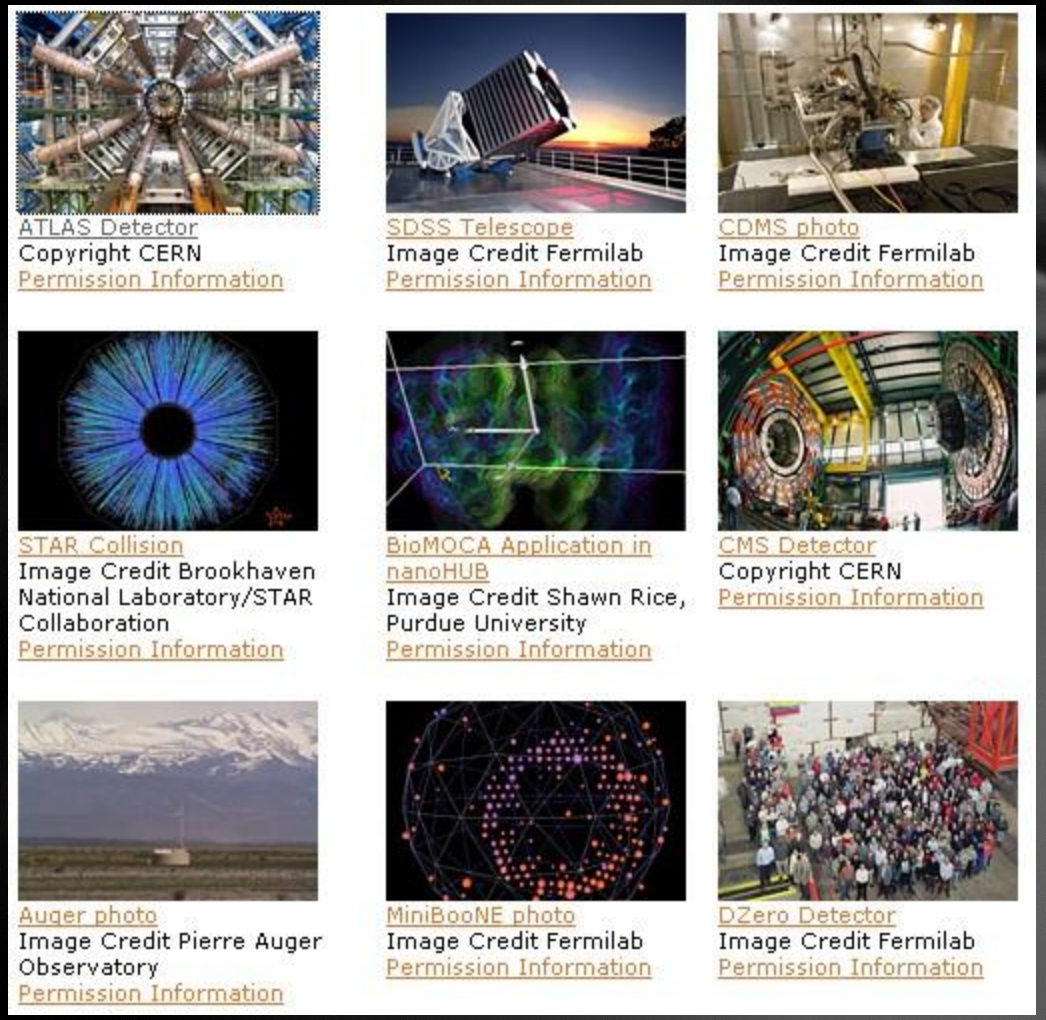

ATLAS Detector
Copyright CERN
Permission Information

SDSS Telescope
Image Credit Fermilab
Permission Information

CDMS photo
Image Credit Fermilab
Permission Information

STAR Collision
Image Credit Brookhaven National Laboratory/STAR Collaboration
Permission Information

BioMOCA Application in nanoHUB
Image Credit Shawn Rice, Purdue University
Permission Information

CMS Detector
Copyright CERN
Permission Information

Auger photo
Image Credit Pierre Auger Observatory
Permission Information

MiniBooNE photo
Image Credit Fermilab
Permission Information

DZero Detector
Image Credit Fermilab
Permission Information

Galactic Plane Workflow

A multiwavelength infrared image atlas of the galactic plane, composed of images at 17 different wavelengths from 1 μm to 70 μm, processed so that they appear to have been measured with a single instrument observing all 17 wavelengths

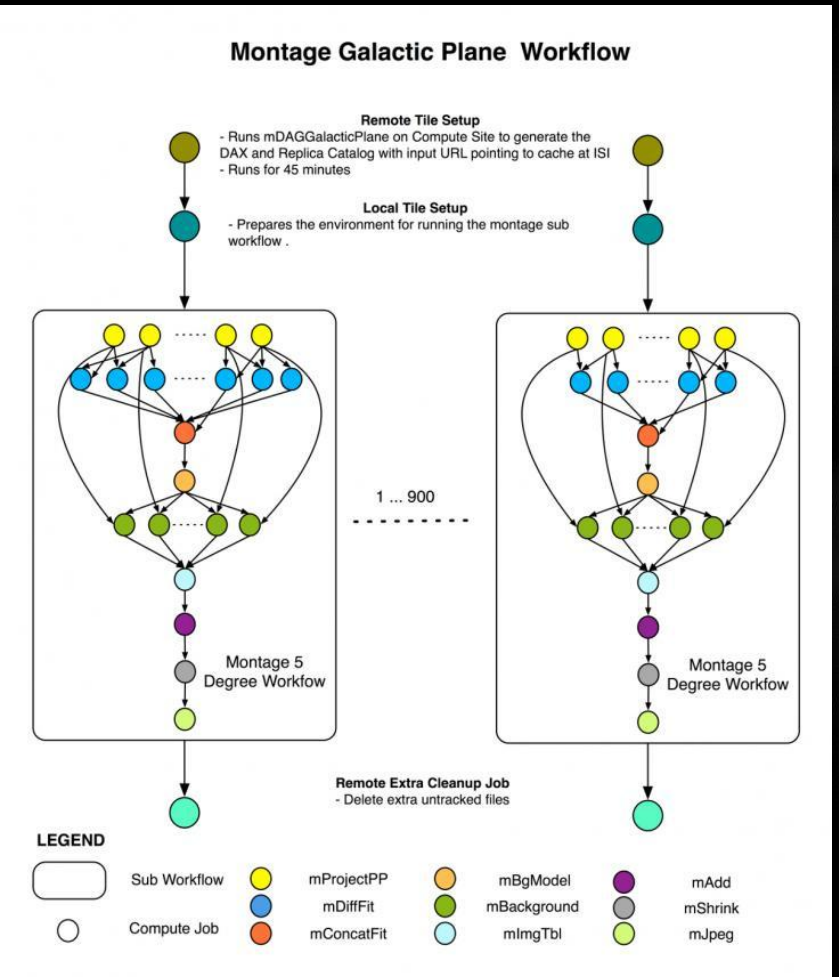

- 360° x 40 ° coverage
- 18 million input files
- 86 TB output dataset
- 17 workflows, each one with 900 sub workflows

Survey	Wavelengths (μm)
Spitzer Space Telescope	
GLIMPSE I, II and 3D	3.6, 4.5, 5.8. 8.0
MIPSGAL I, II	24, 70
All Sky Surveys	
2MASS	1.2, 1.6, 2.2
MSX	8.8, 12.1,14.6, 21.3
WISE †	3.4, 4.6, 12, 22
† Galactic Plane data scheduled for release Spring 2012	

The Extreme Science and Engineering Discovery Environment (XSEDE)

- 9 supercomputers, 3 visualization systems, and 9 storage systems provided by 16 partner institutions
- XSEDE resources are allocated through a peer-reviewed process
- Open to any US open science researcher (or collaborators of US researchers) regardless of funding source
- XSEDE resources are provided at NO COST to the end user through NSF funding (~$100M/year).

Slide credit: Matthew Vaughn

Clouds

Run your own custom virtual machines

But what is provided? What is missing?

Science Clouds

FutureGrid

Commercial Clouds

Amazon EC2, Google Compute, RackSpace

	Linux/UNIX Usage
Standard On-Demand Instances	
Small (Default)	$0.080 per Hour
Medium	$0.160 per Hour
Large	$0.320 per Hour
Extra Large	$0.640 per Hour
Micro On-Demand Instances	
Micro	$0.020 per Hour
High-Memory On-Demand Instances	
Extra Large	$0.450 per Hour
Double Extra Large	$0.900 per Hour
Quadruple Extra Large	$1.800 per Hour
High-CPU On-Demand Instances	
Medium	$0.165 per Hour
Extra Large	$0.660 per Hour
Cluster Compute Instances	
Quadruple Extra Large	$1.300 per Hour
Eight Extra Large	$2.400 per Hour
Cluster GPU Instances	
Quadruple Extra Large	$2.100 per Hour
High-I/O On-Demand Instances	
Quadruple Extra Large	$3.100 per Hour

Data Transfer IN	
All data transfer in	$0.000 per GB
Data Transfer OUT	
First 1 GB / month	$0.000 per GB
Up to 10 TB / month	$0.120 per GB
Next 40 TB / month	$0.090 per GB
Next 100 TB / month	$0.070 per GB
Next 350 TB / month	$0.050 per GB

The application of cloud computing to astronomy: A study of cost and performance
Berriman et.al.

SUMMARY OF PERIODOGRAM CALCULATIONS ON THE AMAZON EC2 CLOUD

		Result
Runtimes	Tasks	631,992
	Mean Task Runtime	6.34 sec
	Jobs	25,401
	Mean Job Runtime	2.62 min
	Total CPU Time	1,113 hr
	Total Wall Time	26.8 hr
Inputs	Input Files	210,664
	Mean Input Size	0.084 MB
	Total Input Size	17.3 GB
Outputs	Output Files	1,263,984
	Mean Output Size	0.124 MB
	Total Output Size	76.52 GB
Cost	Compute Cost	$291.58
	Transfer Cost	$11.48
	Total Cost	**$303.06**

Cloud – what do I have to provide?

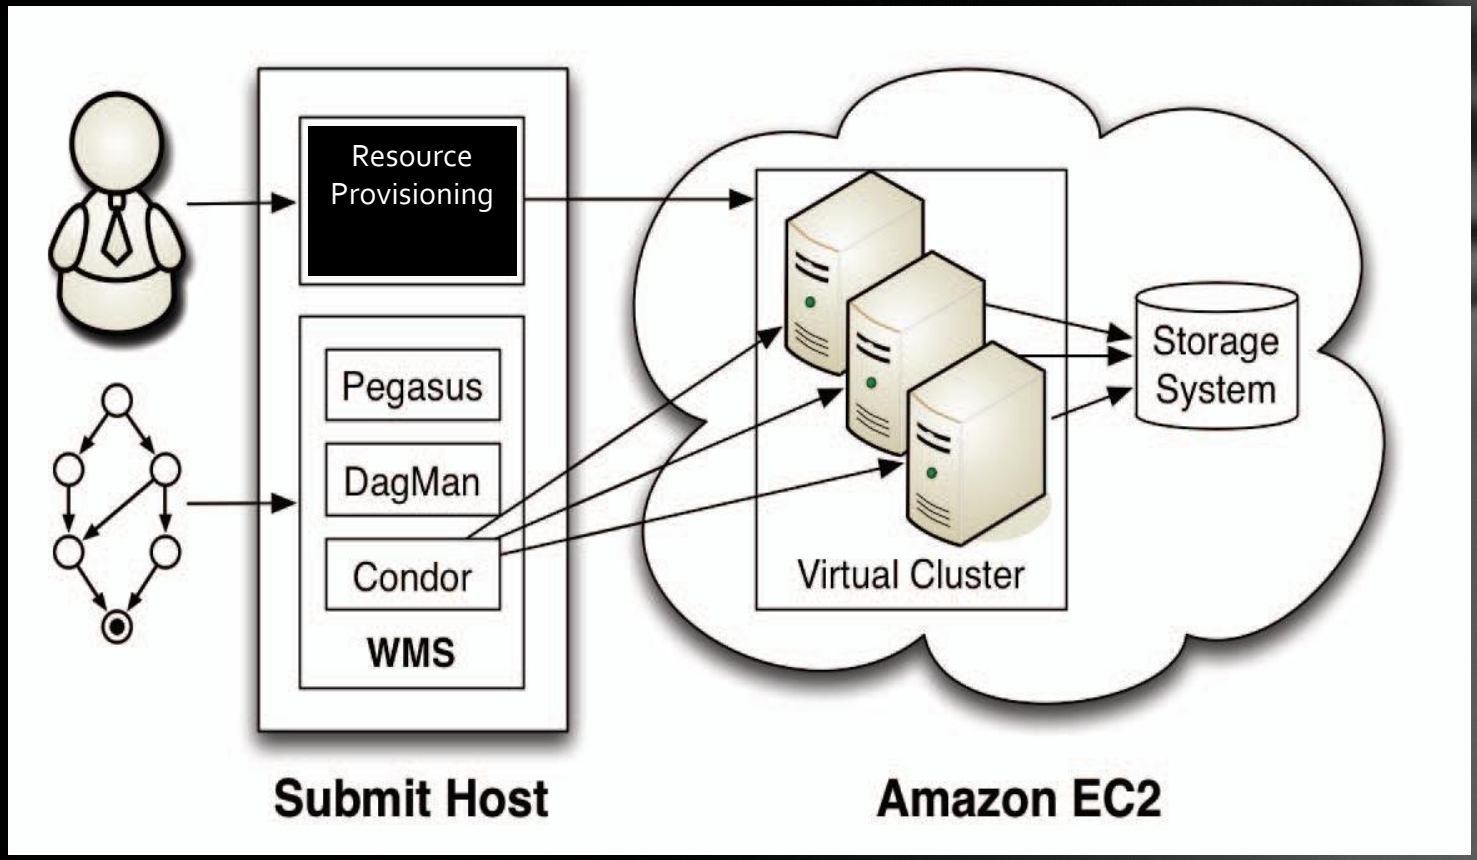

Cloud – Tutorial

This tutorial will take you through the steps of launching the Pegasus Tutorial VM on Amazon EC2 and running a simple workflow. This tutorial is intended for new users who want so get a quick overview of Pegasus concepts and usage. A preconfigured virtual machine is provided so that minimal software installation is required. The tutorial covers the process of starting the VM and of creating, planning, submitting, monitoring, debugging, and generating statistics for a simple workflow.

Oregon datacenter

Image: ami-8643ccb6

http://pegasus.isi.edu/wms/docs/tutorial/

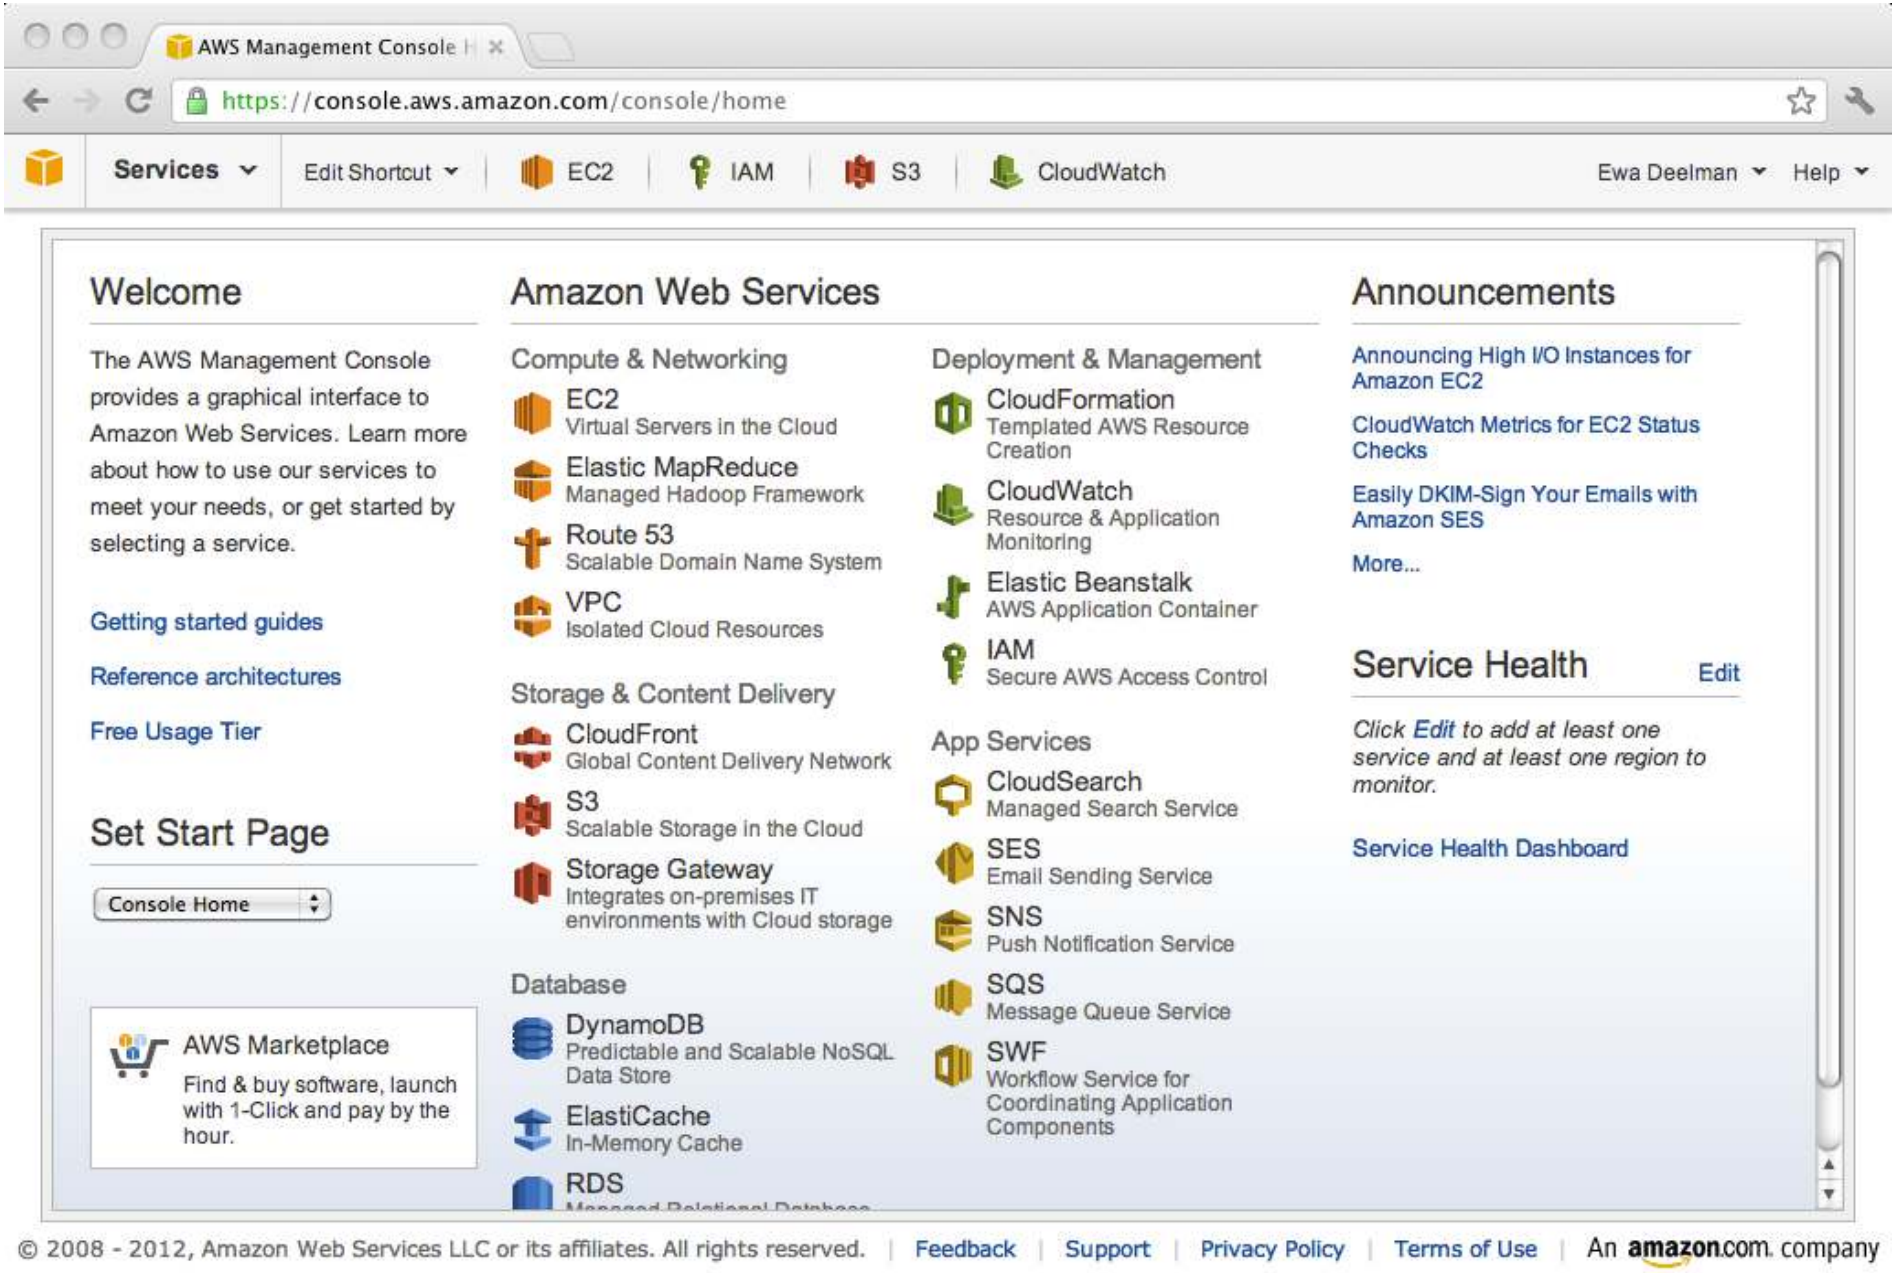
AWS Management Console
https://console.aws.amazon.com/console/home
Services
Edit Shortcut
EC2
IAM
S3
CloudWatch
Ewa Deelman
Help
Welcome
The AWS Management Console provides a graphical interface to Amazon Web Services. Learn more about how to use our services to meet your needs, or get started by selecting a service.
Getting started guides
Reference architectures
Free Usage Tier
Set Start Page
Console Home
AWS Marketplace
Find & buy software, launch with 1-Click and pay by the hour.
Amazon Web Services
Compute & Networking
EC2
Virtual Servers in the Cloud
Elastic MapReduce
Managed Hadoop Framework
Route 53
Scalable Domain Name System
VPC
Isolated Cloud Resources
Storage & Content Delivery
CloudFront
Global Content Delivery Network
S3
Scalable Storage in the Cloud
Storage Gateway
Integrates on-premises IT environments with Cloud storage
Database
DynamoDB
Predictable and Scalable NoSQL Data Store
ElastiCache
In-Memory Cache
RDS
Deployment & Management
CloudFormation
Templated AWS Resource Creation
CloudWatch
Resource & Application Monitoring
Elastic Beanstalk
AWS Application Container
IAM
Secure AWS Access Control
App Services
CloudSearch
Managed Search Service
SES
Email Sending Service
SNS
Push Notification Service
SQS
Message Queue Service
SWF
Workflow Service for Coordinating Application Components
Announcements
Announcing High I/O Instances for Amazon EC2
CloudWatch Metrics for EC2 Status Checks
Easily DKIM-Sign Your Emails with Amazon SES
More...
Service Health
Edit
Click Edit to add at least one service and at least one region to monitor.
Service Health Dashboard
© 2008 - 2012, Amazon Web Services LLC or its affiliates. All rights reserved.
Feedback
Support
Privacy Policy
Terms of Use
An amazon.com company

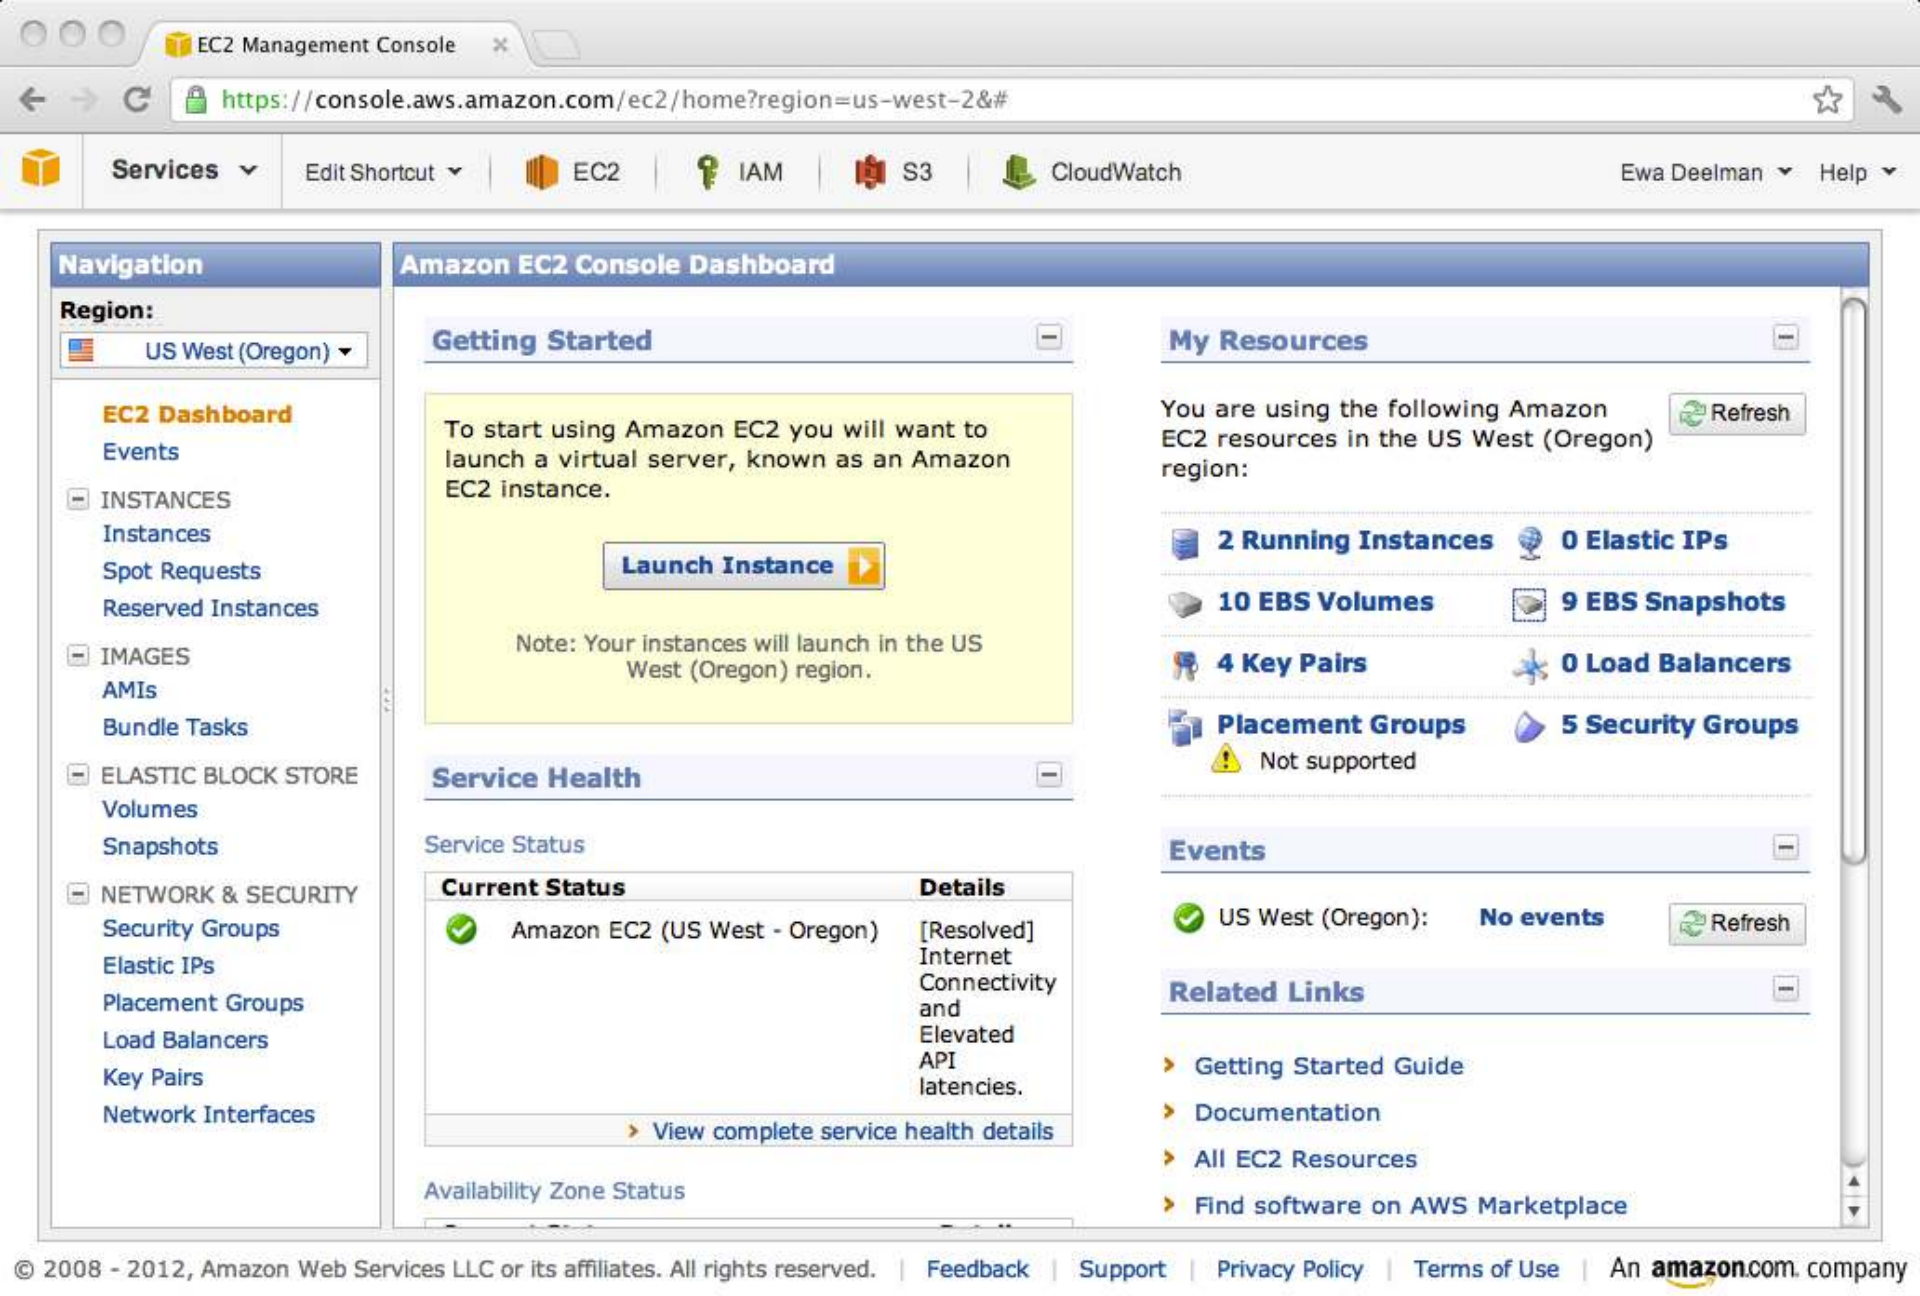
EC2 Management Console
https://console.aws.amazon.com/ec2/home?region=us-west-2&#
Services
Edit Shortcut
EC2
IAM
S3
CloudWatch
Ewa Deelman
Help
Navigation
Region:
US West (Oregon)
EC2 Dashboard
Events
INSTANCES
Instances
Spot Requests
Reserved Instances
IMAGES
AMIs
Bundle Tasks
ELASTIC BLOCK STORE
Volumes
Snapshots
NETWORK & SECURITY
Security Groups
Elastic IPs
Placement Groups
Load Balancers
Key Pairs
Network Interfaces
Amazon EC2 Console Dashboard
Getting Started
To start using Amazon EC2 you will want to launch a virtual server, known as an Amazon EC2 instance.
Launch Instance
Note: Your instances will launch in the US West (Oregon) region.
Service Health
Service Status
Current Status
Details
Amazon EC2 (US West - Oregon)
[Resolved] Internet Connectivity and Elevated API latencies.
View complete service health details
Availability Zone Status
My Resources
You are using the following Amazon EC2 resources in the US West (Oregon) region:
Refresh
2 Running Instances
0 Elastic IPs
10 EBS Volumes
9 EBS Snapshots
4 Key Pairs
0 Load Balancers
Placement Groups
Not supported
5 Security Groups
Events
US West (Oregon):
No events
Refresh
Related Links
Getting Started Guide
Documentation
All EC2 Resources
Find software on AWS Marketplace
© 2008 - 2012, Amazon Web Services LLC or its affiliates. All rights reserved.
Feedback
Support
Privacy Policy
Terms of Use
An amazon.com company

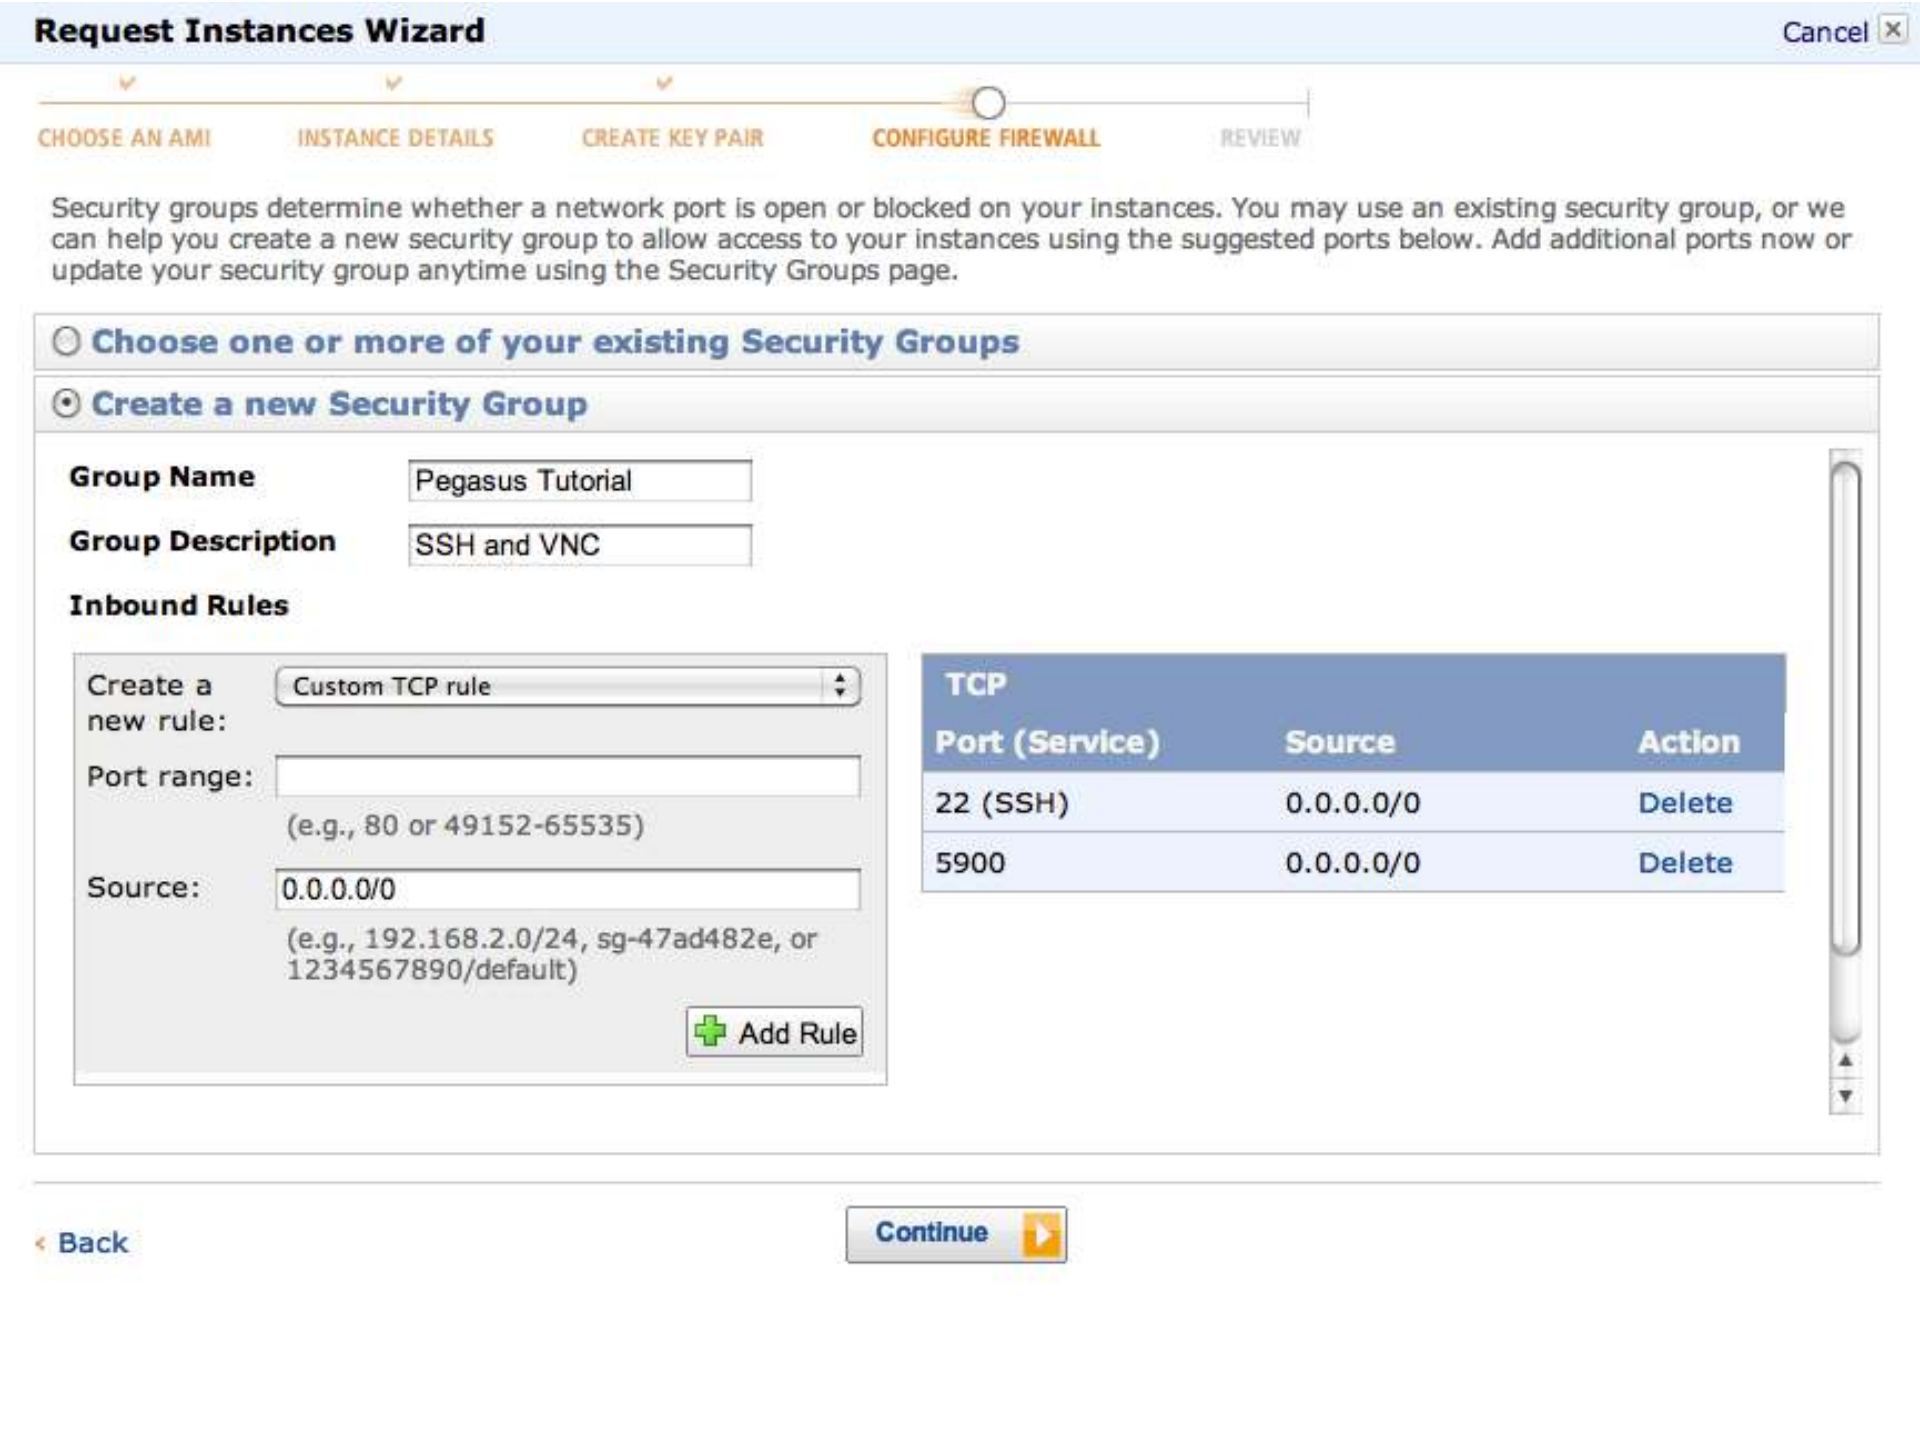
Request Instances Wizard
Cancel
CHOOSE AN AMI
INSTANCE DETAILS
CREATE KEY PAIR
CONFIGURE FIREWALL
REVIEW
Security groups determine whether a network port is open or blocked on your instances. You may use an existing security group, or we can help you create a new security group to allow access to your instances using the suggested ports below. Add additional ports now or update your security group anytime using the Security Groups page.
Choose one or more of your existing Security Groups
Create a new Security Group
Group Name
Pegasus Tutorial
Group Description
SSH and VNC
Inbound Rules
Create a new rule:
Custom TCP rule
Port range:
(e.g., 80 or 49152-65535)
Source:
0.0.0.0/0
(e.g., 192.168.2.0/24, sg-47ad482e, or 1234567890/default)
Add Rule
TCP
Port (Service)
Source
Action
22 (SSH)
0.0.0.0/0
Delete
5900
0.0.0.0/0
Delete
Back
Continue

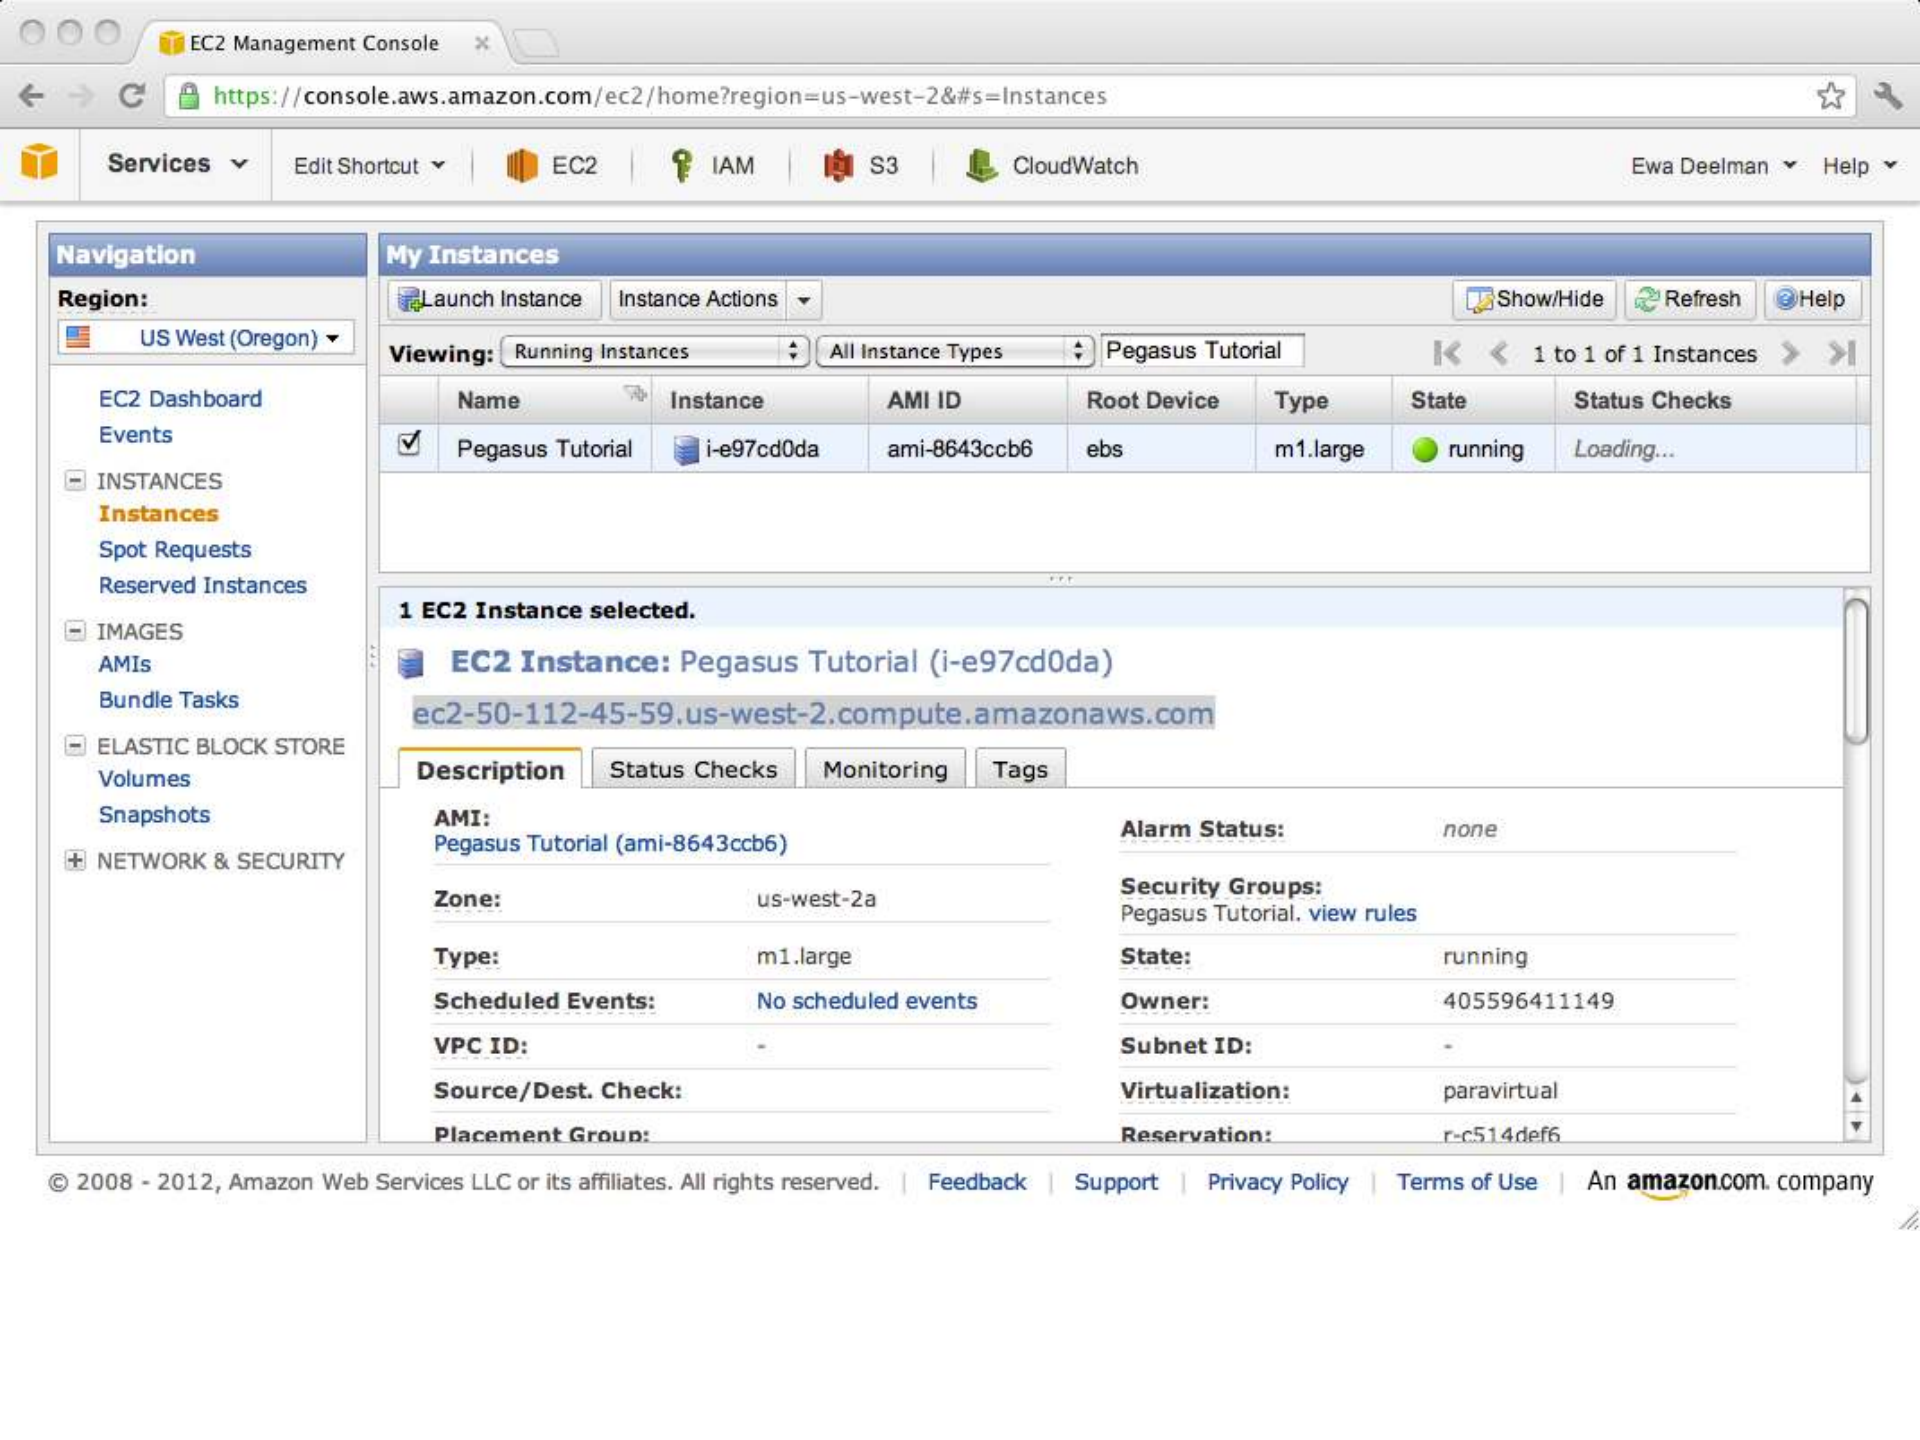
EC2 Management Console
https://console.aws.amazon.com/ec2/home?region=us-west-2&#s=Instances
Services
Edit Shortcut
EC2
IAM
S3
CloudWatch
Ewa Deelman
Help
Navigation
Region:
US West (Oregon)
EC2 Dashboard
Events
INSTANCES
Instances
Spot Requests
Reserved Instances
IMAGES
AMIs
Bundle Tasks
ELASTIC BLOCK STORE
Volumes
Snapshots
NETWORK & SECURITY
My Instances
Launch Instance
Instance Actions
Show/Hide
Refresh
Help
Viewing:
Running Instances
All Instance Types
Pegasus Tutorial
1 to 1 of 1 Instances
Name
Instance
AMI ID
Root Device
Type
State
Status Checks
Pegasus Tutorial
i-e97cd0da
ami-8643ccb6
ebs
m1.large
running
Loading...
1 EC2 Instance selected.
EC2 Instance: Pegasus Tutorial (i-e97cd0da)
ec2-50-112-45-59.us-west-2.compute.amazonaws.com
Description
Status Checks
Monitoring
Tags
AMI:
Pegasus Tutorial (ami-8643ccb6)
Zone:
us-west-2a
Type:
m1.large
Scheduled Events:
No scheduled events
VPC ID:
-
Source/Dest. Check:
Placement Group:
Alarm Status:
none
Security Groups:
Pegasus Tutorial. view rules
State:
running
Owner:
405596411149
Subnet ID:
-
Virtualization:
paravirtual
Reservation:
r-c514def6
© 2008 - 2012, Amazon Web Services LLC or its affiliates. All rights reserved.
Feedback
Support
Privacy Policy
Terms of Use
An amazon.com company

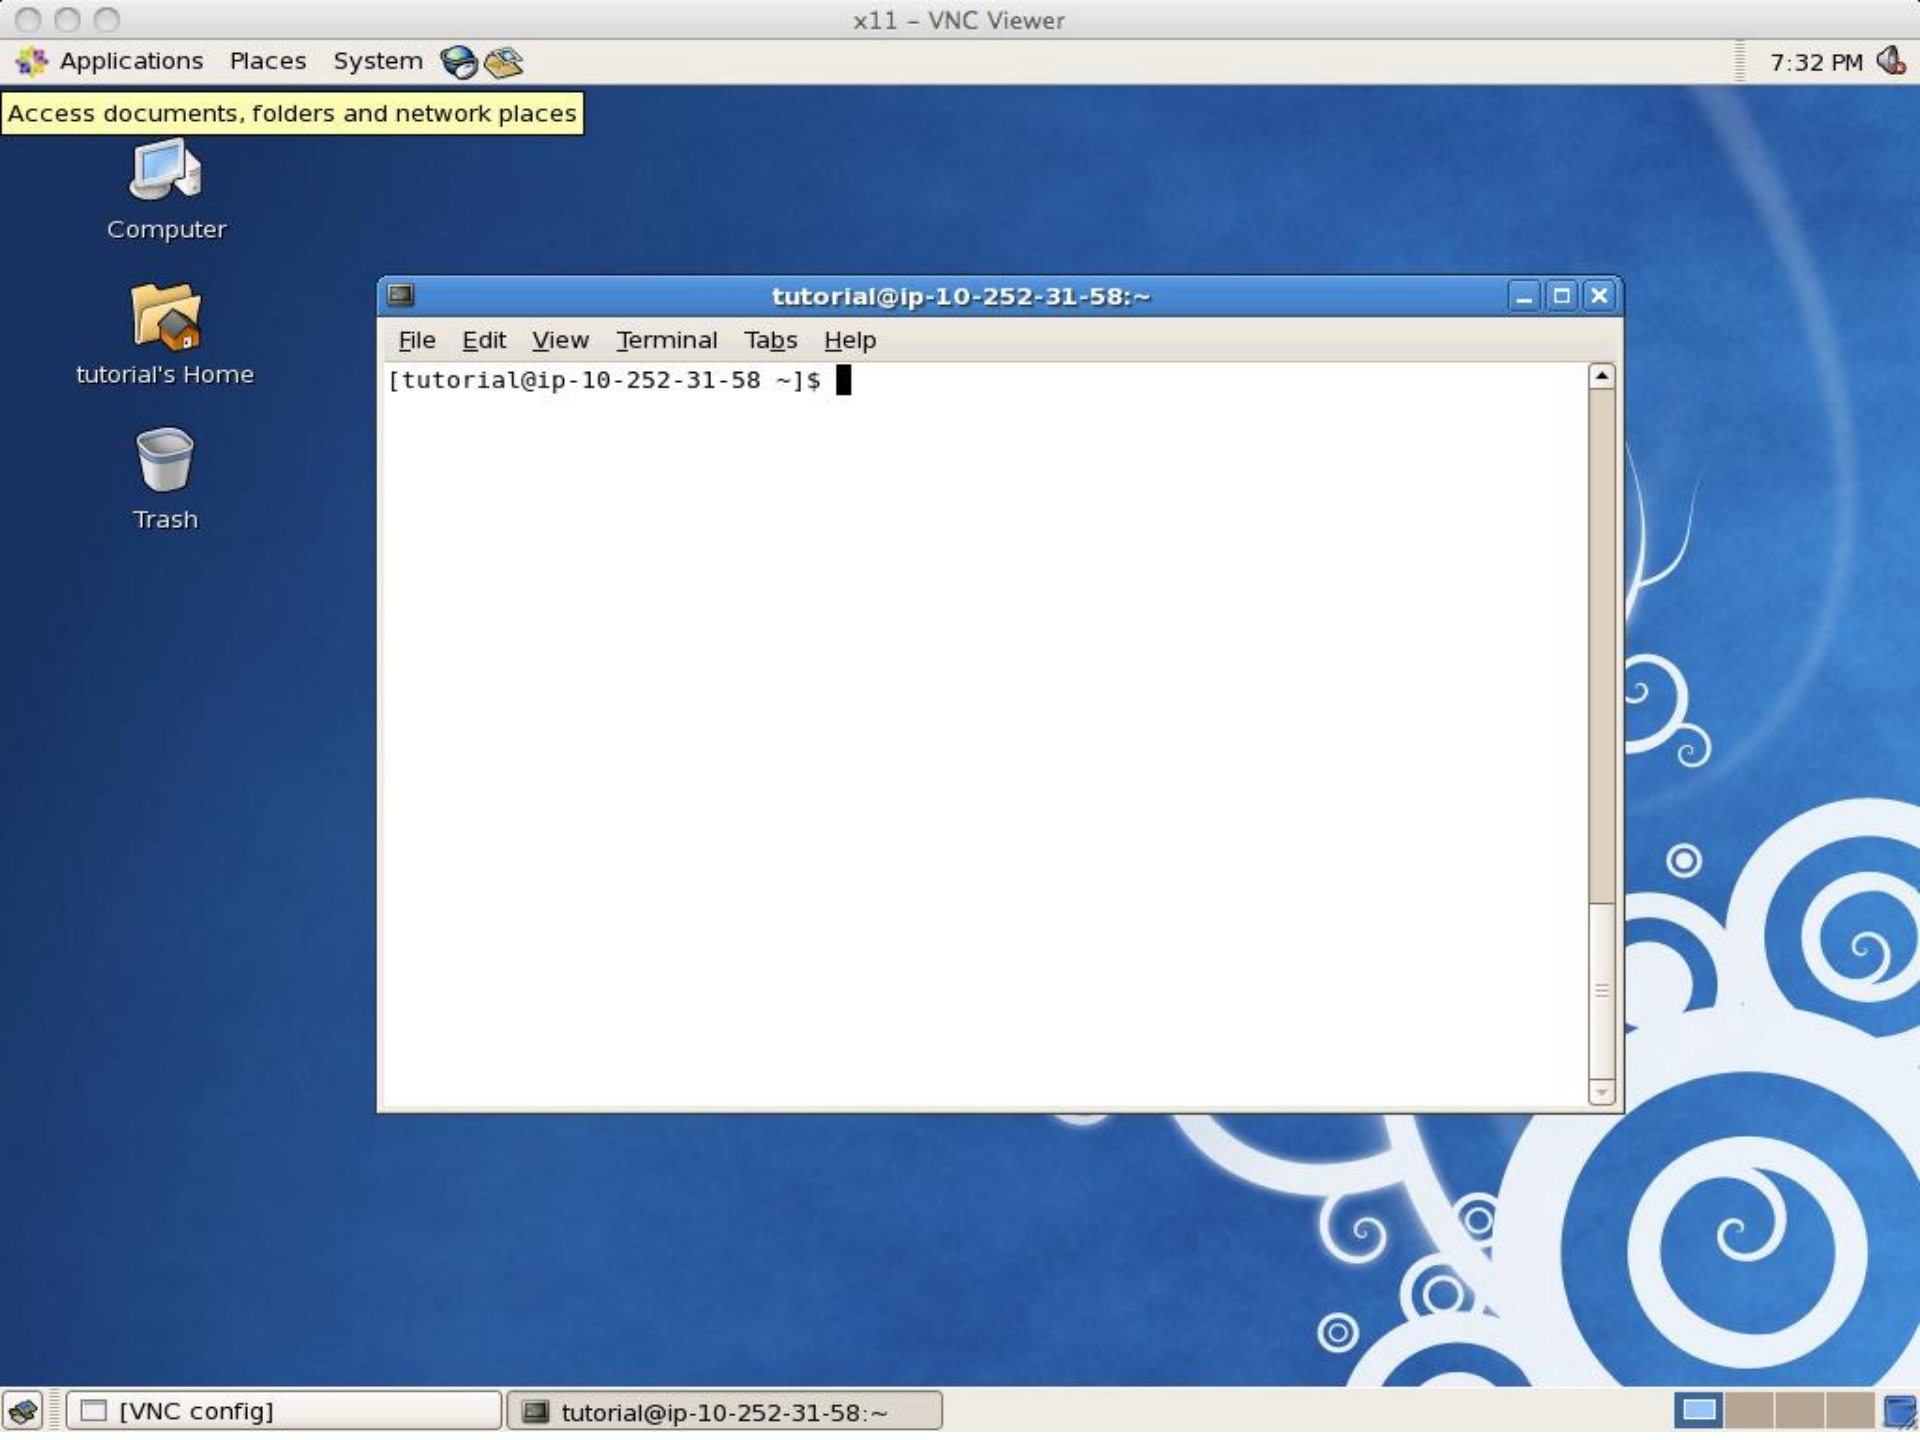
x11 – VNC Viewer
Applications
Places
System
7:32 PM
Access documents, folders and network places
Computer
tutorial's Home
Trash
tutorial@ip-10-252-31-58:~
File Edit View Terminal Tabs Help
[tutorial@ip-10-252-31-58 ~]$
[VNC config]
tutorial@ip-10-252-31-58:~

Thank you!

Pegasus: *http://pegasus.isi.edu*

pegasus-support@isi.edu