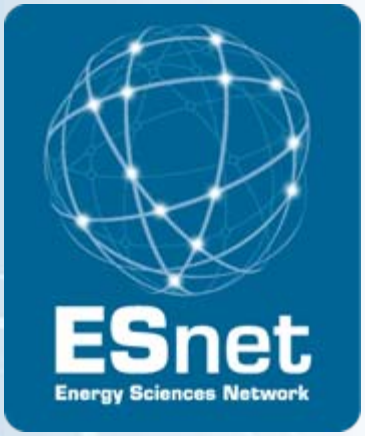

Implications of ESnet Site Reliance on Cloud Services

Greg Bell, ESnet

ESCC – July 14, 2010

Topics

- recent trends and developments
- summary of Phil's ESCC survey
- questions and open discussion

A Simple Definition of ‘Cloud’ from SC09

‘Cloud’ implies an architecture in which computing systems come in two sizes, tiny and huge.

- courtesy of Urs Hoelzle (SVP at Google), SC09 Masterworks Talk ‘Warehouse Scale Computers’
- tiny: smart phones, net/notebooks, thin clients, tablets
- huge: ‘warehouse-scale computers’: ~$500M, ~50MW
  - your data & apps live here
- optimized for efficiency at either scale
  - battery constraints
  - OpEx @ 50MW

Beyond *'Storm Cloud vs. Fluffy Cloud'*

Recent Trends

Some ESnet sites – by no means all of them – actively cloud-sourcing.

- collaborative services
- business apps
- backup services
- app hosting
- CPU / disk rental
  - HPC
  - genomics workflow
  - CS research
  - enterprise IT

News from the Past Two Days

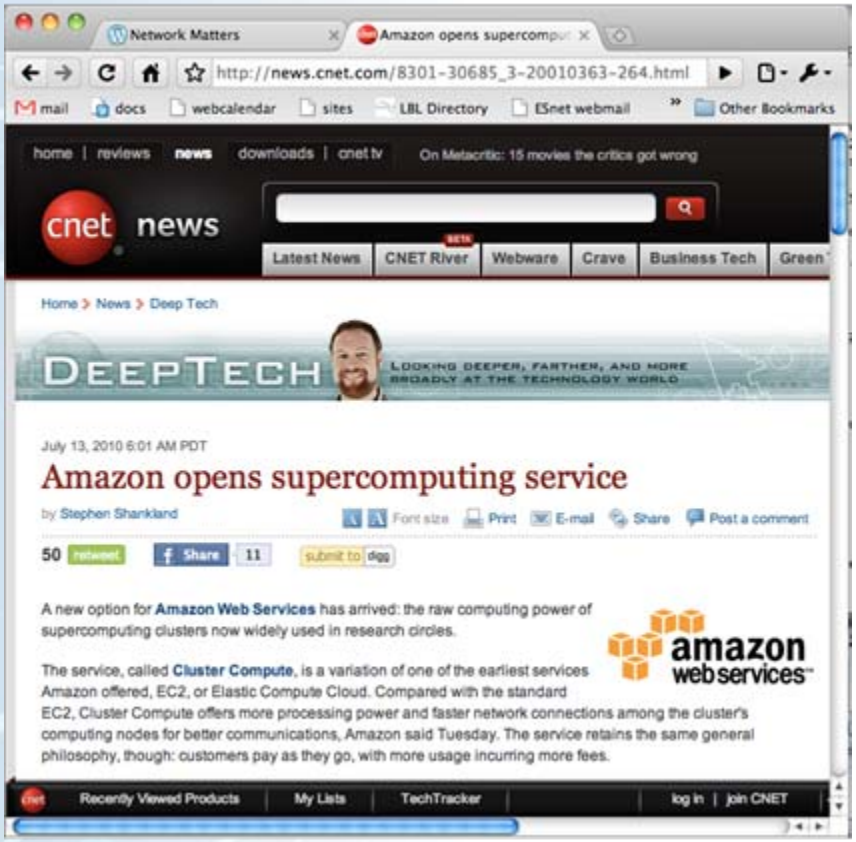

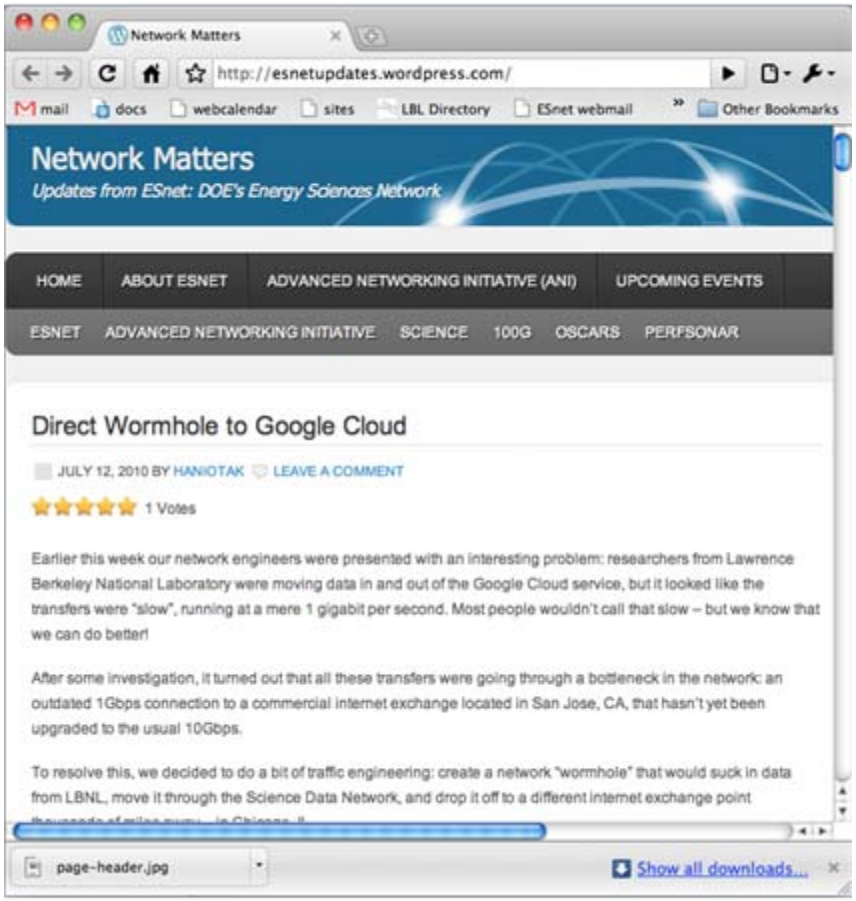

Reviewing ESCC Survey Results

1) Does your site currently [or plan to] depend on cloud-based / remotely-hosted critical services?

- Yes, this is a strategic direction [27.3%, 3/11]
- Some (limited) cloud-based dependencies [0%, 0/11]
- Under consideration, but no definite decision [72.7%, 8/11]
- Not a strategic direction at this time [0%, 0/11]

Reviewing ESCC Survey Results

2) What is your level of reliance on ESnet to access these services?

- Completely reliant on ESnet [45.5%, 5/11]
- Other network paths available [36.4%, 4/11]
- Contingency plan(s) mitigate impact of access problem [9.1%, 1/11]
- Not applicable; no dependencies on remote services [9.1%, 1/11]

Reviewing ESCC Survey Results

3) Do you anticipate significant additional resource demands (b/w capacity, accessibility) on ESnet to meet your service needs?

- **Yes [36.4%, 4/11]**
- No [18.2%, 2/11]
- Unclear [45.5%, 5/11]
- Not applicable; not remote dependencies [0%, 0/11]

Reviewing ESCC Survey Results

4) What other question or concern relative to this issue would you like to see brought up at the meeting?

"Reliable access to services."

"Our biggest issue to moving more to the cloud is physical network diversity and SLA maturity (on our part) with the service provider."

"We don't currently have plans to rely on cloud services that would cause a significant increase in demand (next ~1 year), but it's possible that could change further out."

"Traffic volume could increase significantly but stay with the capacity of existing circuits connecting our site to ESnet."

"Redundant connection to ORNL, redundant equipment at our site."

Possible Discussion Questions

For sites anticipating ‘significant additional resource demands’ on ESnet…

- what needs are you most focused on (connection b/w, peering b/w, reduced latency, dedicated circuits, improved resiliency, etc)?
- what specific scenarios do you anticipate?

Does your site have any research collaborations with cloud-service providers?

Will major cloud-sourcing projects kick off in the next year?

Have you discovered any surprising dependencies on remote services?

Are you planning to source backup services?

Are there any questions we’re not asking, but should be?

Thank You

- feedback or questions? grbell@es.net