

MONITORING AND CONTROLLING THE SMART ELECTRIC POWER GRID.

March 30, 2012

Ken Birman, Cornell University (ken@cs.cornell.edu)



Click Here!

Advantage: Cloud



3

- The distributed computing technology landscape is being transformed by *cloud computing*
 - ▣ Defn: ***the provision of computational resources on demand via a network (Wikipedia)***
 - ▣ Often uses a shared infrastructure, virtualization
 - ▣ Potential to amass huge amounts of data in a single data center, making it possible to replace *model-driven decision making* with *data-driven approaches*
- Cloud computing has many economic and scalability advantages



Public clouds and private clouds

4

- Many assume that power grid could never run on cloud
 - ▣ They envision some form of dedicated, special infrastructure
 - ▣ But in technology, it gets very expensive to roll your own
 - ▣ Cloud story is one of leveraging economies of scale

- Today:
 - ▣ Why is cloud computing winning?
 - ▣ How good is the fit for future electric power applications?
 - ▣ How can we close any gaps we identify?



The “way” of cloud computing

5

- Cloud computing platforms have a characteristic style
 - **Client system** is a browser or an application “pretending” to be a browser (web services). If a request times out the client will retry it
 - **Cloud platform has a massive number of servers facing the clients.** These are “stateless” and handle client requests with a focus on speed. If one fails, there is no warning. The cluster manager just restarts it, perhaps elsewhere
 - **Behind the servers is a row of stateful (usually database) servers.** These operate off some form of work “queues”, usually implemented as enterprise service buses (ESBs). Think “publish-subscribe” with logging. They push updates to the client-facing systems.
 - **Finally, there is a big pool of machines to run “backend” applications.** This looks like a loosely coupled collection of Linux machines sharing a file system and a lock manager, and often programmed mostly with Hadoop (MapReduce).

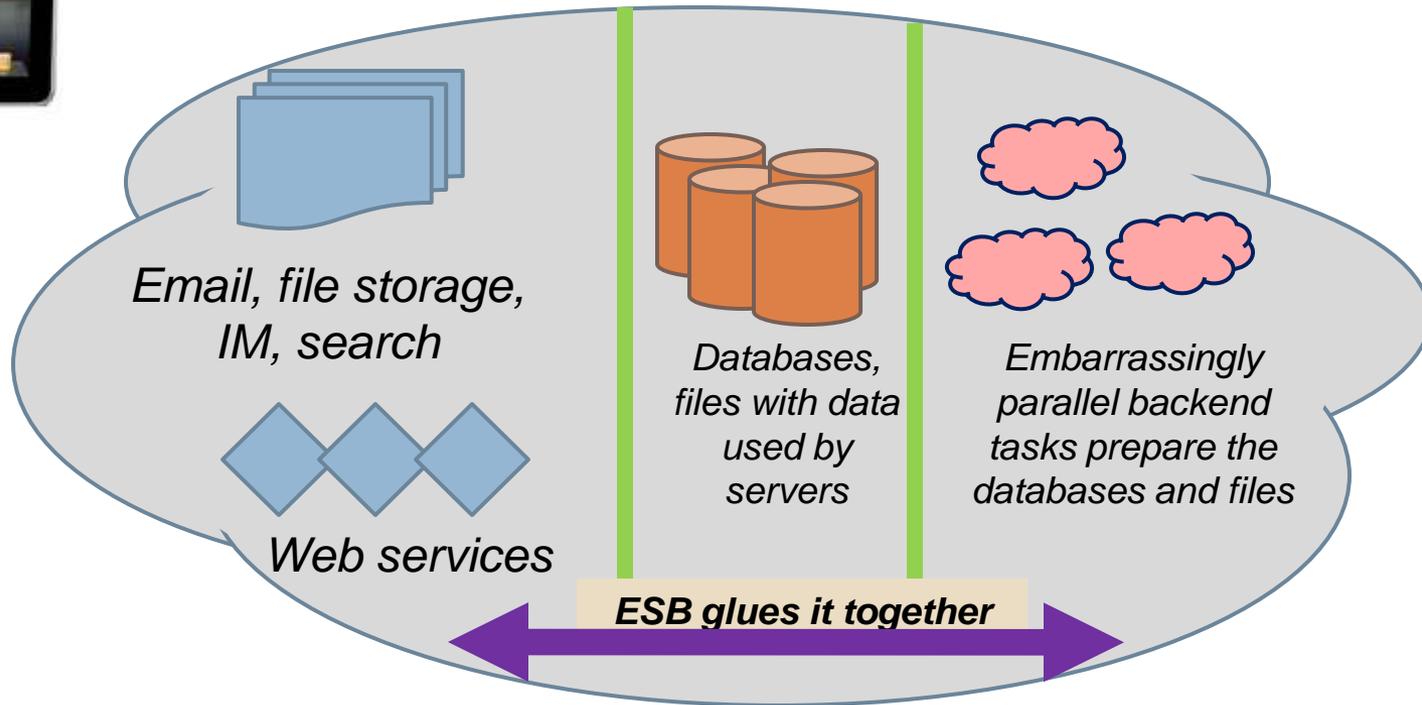


The “way” of cloud computing

6



Client systems use web technologies or run “apps”



Google/IBM/Amazon/Facebook host the services
Computational Challenges for Wind Prediction March 30, 2012



Clouds are cheaper... much cheaper...

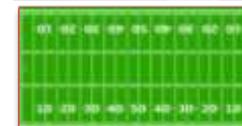
7

Range in size from “edge” facilities to megascale.

Incredible economies of scale

Approximate costs for a small size center (1K servers) and a larger, 50K server center.

Technology	Cost in small-sized Data Center	Cost in Large Data Center	Cloud Advantage
Network	\$95 per Mbps/month	\$13 per Mbps/month	7.1
Storage	\$2.20 per GB/month	\$0.40 per GB/month	5.7
Administration	~140 servers/Administrator	>1000 Servers/Administrator	7.1



Each data center is
11.5 times
the size of a football field



Elasticity... Weak Consistency

8

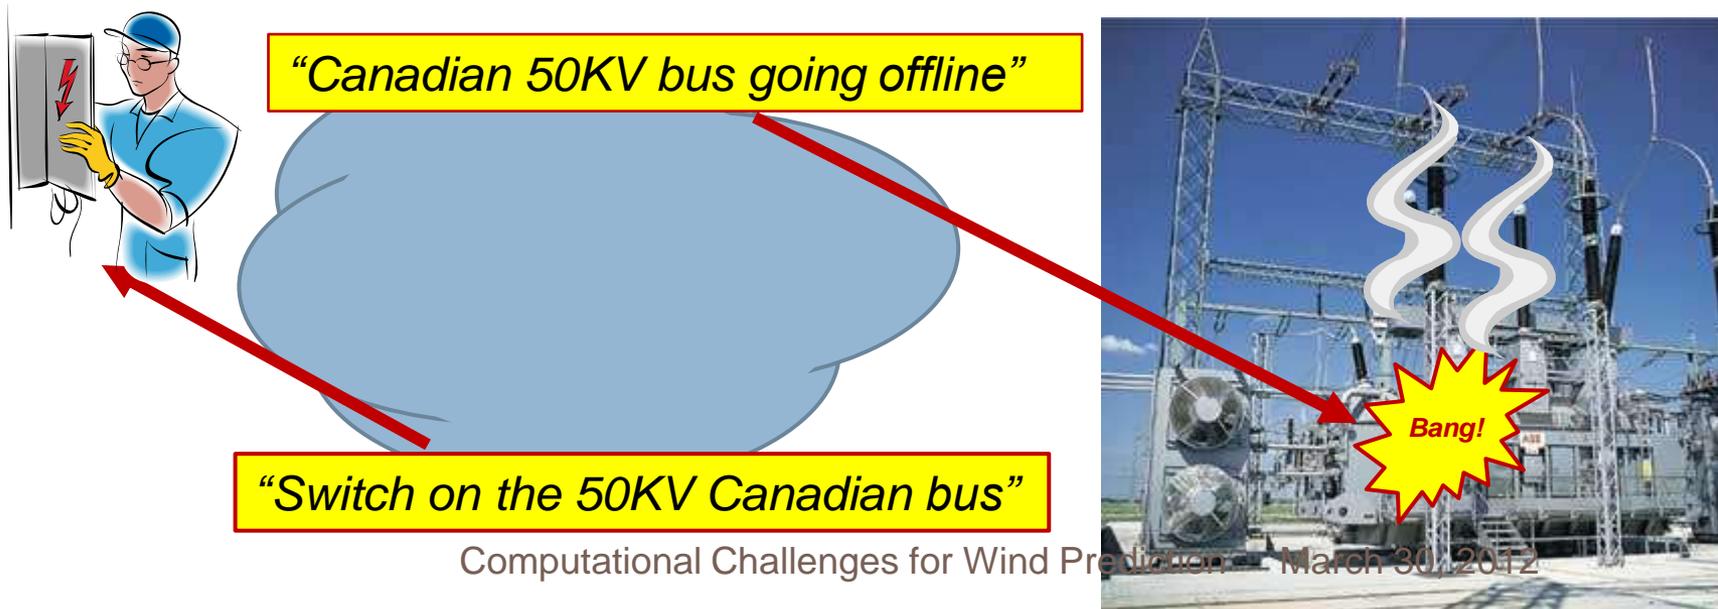
- “Spin on a dime” to accommodate changing load
 - ▣ App style departs from classic client-server computing: weak consistency, few “guarantees”
 - ▣ Instead, they use the “BASE” approach (Basically Available, Soft State, Eventual Consistency).
 - ▣ BASE reduces cost of cloud service development and facilitates elasticity but weakens consistency and fault-tolerance to achieve better scaling
- Can expose client systems to inconsistency



Dangers of Inconsistency

9

- Inconsistency causes bugs
 - ▣ Cloud control system speaks with “two voices”
 - ▣ In physical infrastructure settings, consequences can be very costly



Our Study of this topic

10

- Commissioned by LBL in 2010, results in a white paper

- Our team looked at various use scenarios
 - ▣ National-scale phasor monitoring: huge data rates
 - ▣ Smart appliances that the grid “talks to” to schedule load cycles: vast numbers of clients
 - ▣ Future SCADA in a disaggregated grid: new styles of control optimized for data-rich environments

- In each case we conclude that the scale of the need dwarfs prior styles of computing and that only a cloud model makes sense today



... We also find a serious mismatch

11

- Several aspects of the cloud model are mismatched to the high assurance needs of the future smart power grid
 - ▣ The Internet, which links the client to the cloud, is poorly secured and sometimes unavailable
 - ▣ Cloud application software, in the current generation of systems, embraces inconsistency, yet many power control applications need strong properties to be operated safely
 - ▣ Cloud security limited to a very narrow policy



Executive Summary of Findings

12

Support for scalable real-time services. Today's cloud is fast enough, but sometimes gains speed by cutting corners in ways that control applications can't tolerate.

Support for scalable, consistency guaranteed, fault-tolerant services. Control of a physical infrastructure often brings strong consistency or other assurance needs, for which the cloud currently lacks technology.

Protection of Private Data. Current cloud platforms do a poor job of protecting private data, because privacy is at odds with the cloud's economic incentives.

Highly Assured Internet Routing. In today's Internet, consumers often experience brief periods of loss of connectivity. Redundant routing could overcome this issue.



Summary of Highest Priority Research Topics

13

- ❑ **Quantify the kinds of guarantees that cloud computing solutions can offer.**
- ❑ **Quantify the kinds of guarantees required for smart grid control: match paradigm to feasible guarantees**
- ❑ **Reintroduce strong trust properties in cloud settings.**
- ❑ **Better quantify the possible attacks against a computer-controlled smart grid.**
- ❑ **Build an Internet with better availability properties.**
- ❑ **Improve attack tolerance.**



With ARPAAe support from GENI...

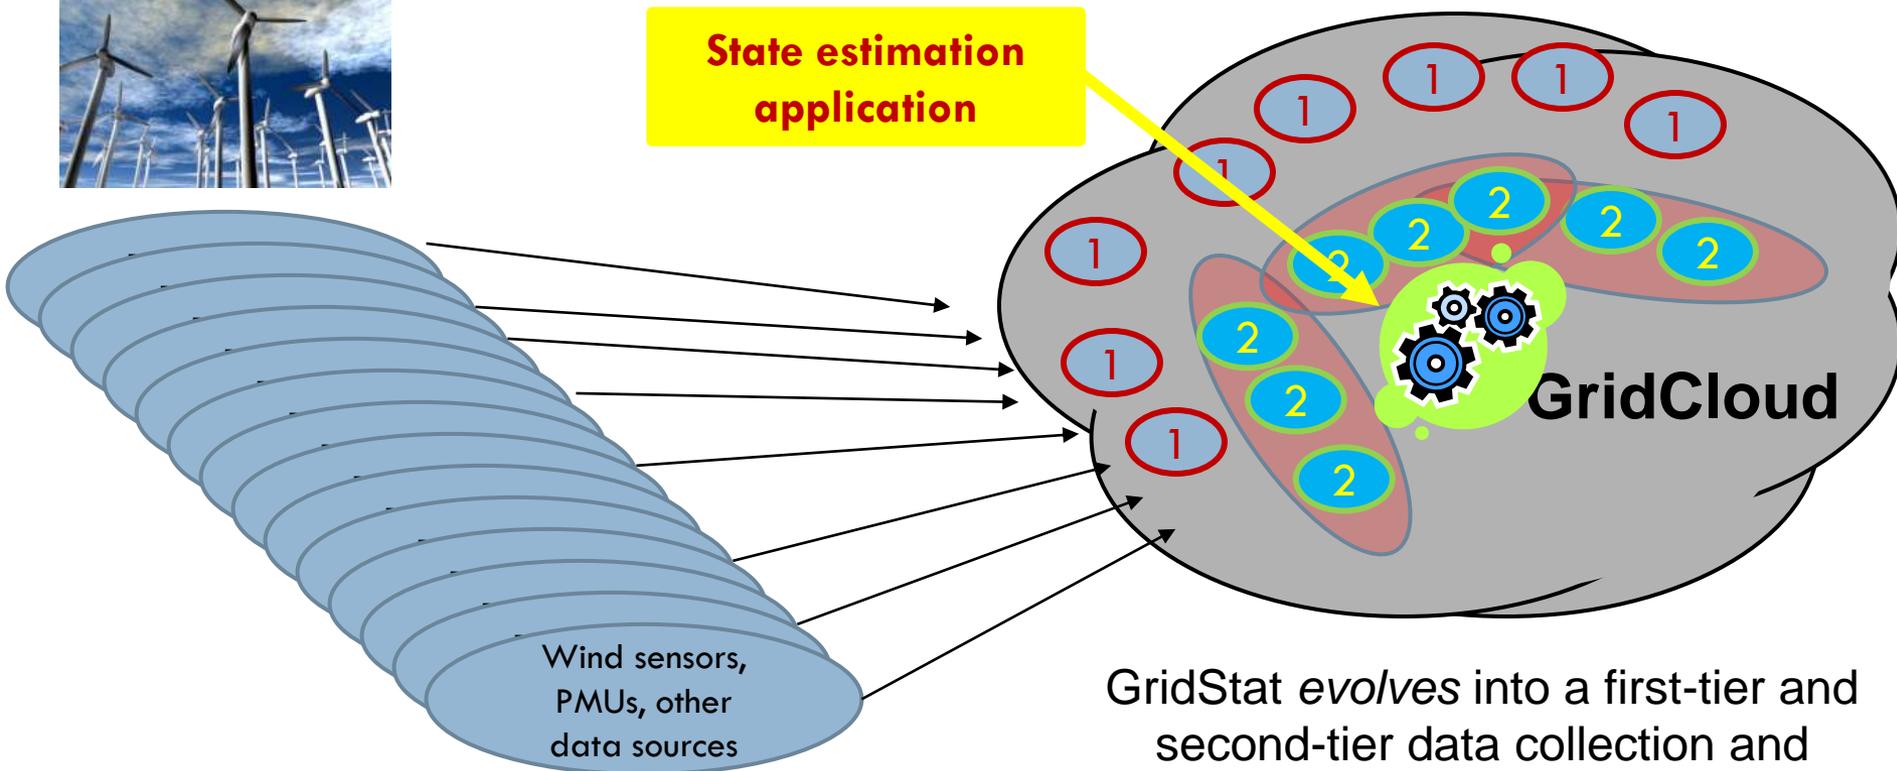
14

- We're exploring this topic
- Combining the Washington State University GridStat and GridSim infrastructures with Cornell's Isis² high assurance cloud computing platform
- Goal: a new kind of “operating system” for high assurance cloud computing



Coming Soon: GridCloud

15

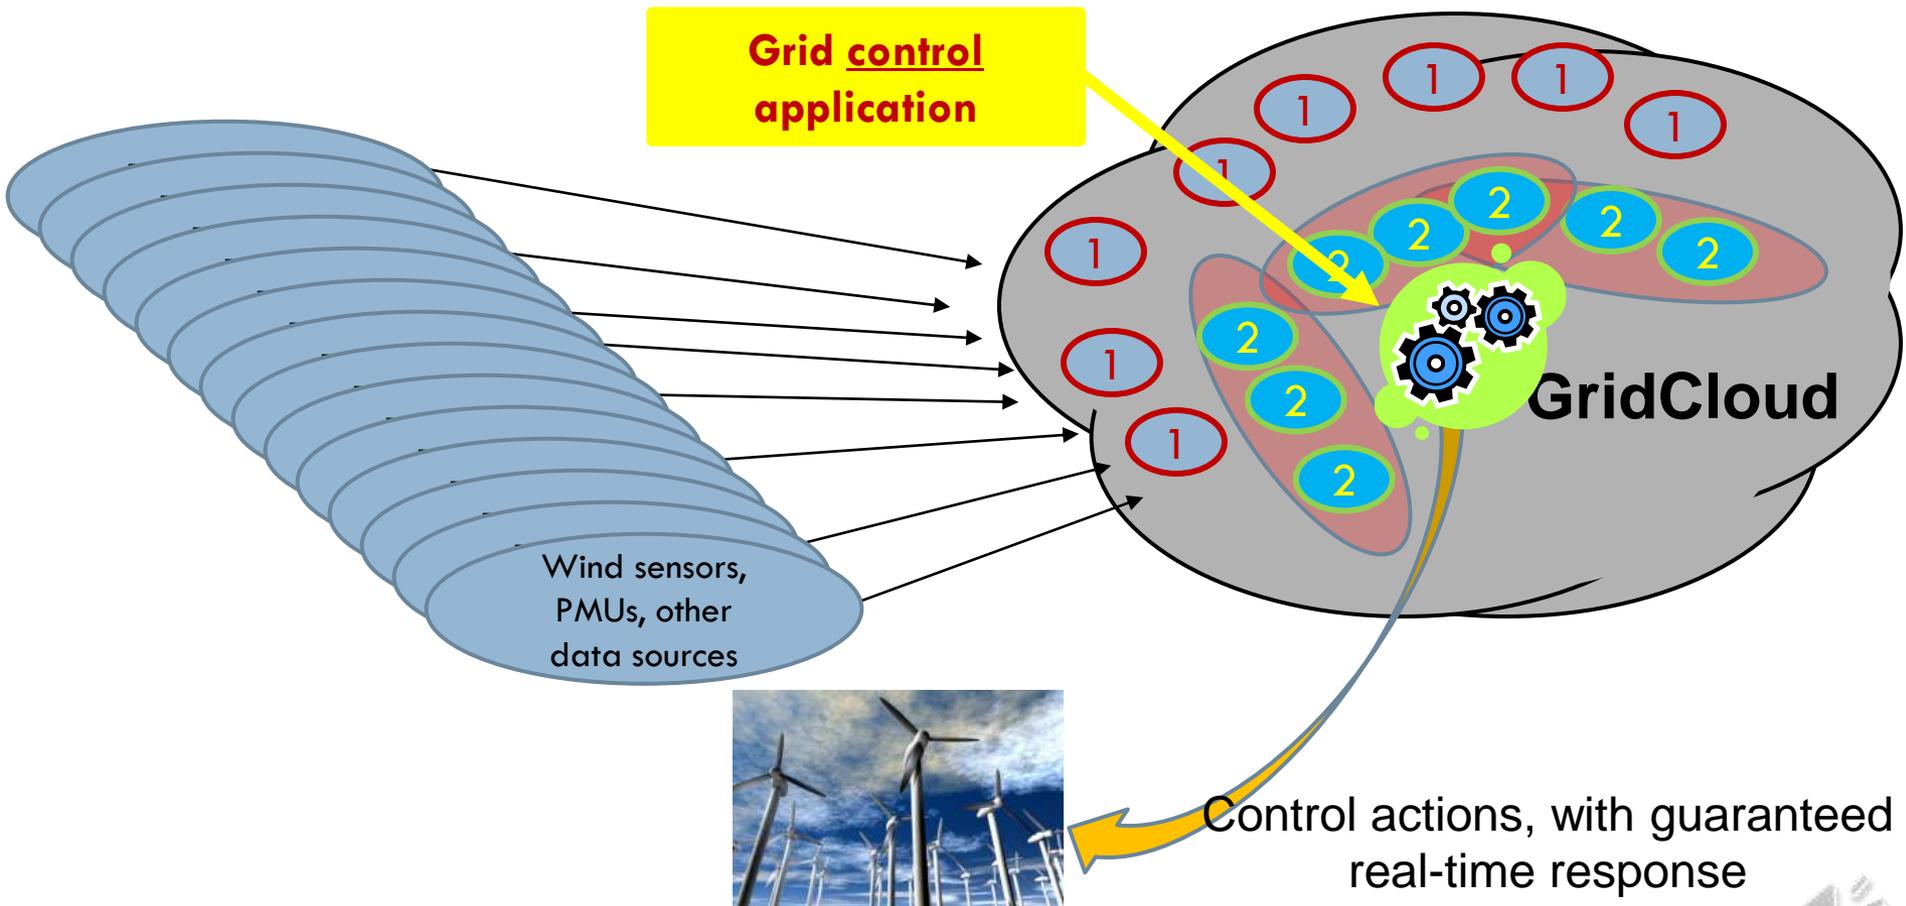


GridStat *evolves* into a first-tier and second-tier data collection and management infrastructure that can host control applications



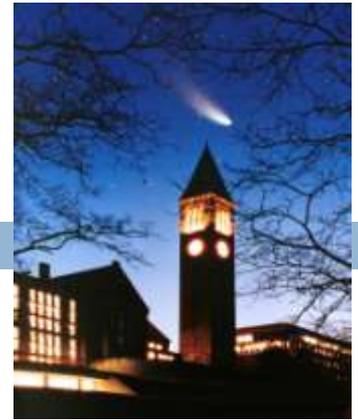
Eventually, close the loop

16



Summary?

17



- The word on the street is that cloud computing will rule but that the cloud can't support high assurance control apps

- The Cornell/WSU team doesn't accept that limitation
 - ▣ We'll build GridCloud to prove our point
 - ▣ But we'll need to replace our "demo" applications with real ones to actually deploy GridControl in real smart-grid scenarios, and for that, we need you!

