A Market Approach for Handling Power Emergencies in Multi-Tenant Data Center

Mohammad A. Islam, Xiaoqi Ren, Shaolei Ren, Adam Wierman, and Xiaorui Wang

What makes up the costs in data centers?

Amortized Cost	Component	Sub-Components
∼45%	Servers	CPU, memory, storage systems
∼25%	Infrastructure	Power distribution and cooling
∼15%	Power draw	Electrical utility costs
∼15%	Network	Links, transit, equipment

Source: A. Greenberg, J. Hamilton, D. A. Maltz, and P. Patel. 2008. The cost of a cloud: research problems in data center networks. SIGCOMM Comput. Commun. Rev.

What makes up the costs in data centers?

Capital Expenditure (CapEx)

CapEx > 1.5×OpEx !

Operational Expenditure (OpEx)

Amortized Cost	Component	Sub-Components
∼25%	Infrastructure	Power distribution and cooling
∼15%	Power draw	Electrical utility costs

Source: A. Greenberg, J. Hamilton, D. A. Maltz, and P. Patel. 2008. The cost of a cloud: research problems in data center networks. SIGCOMM Comput. Commun. Rev.

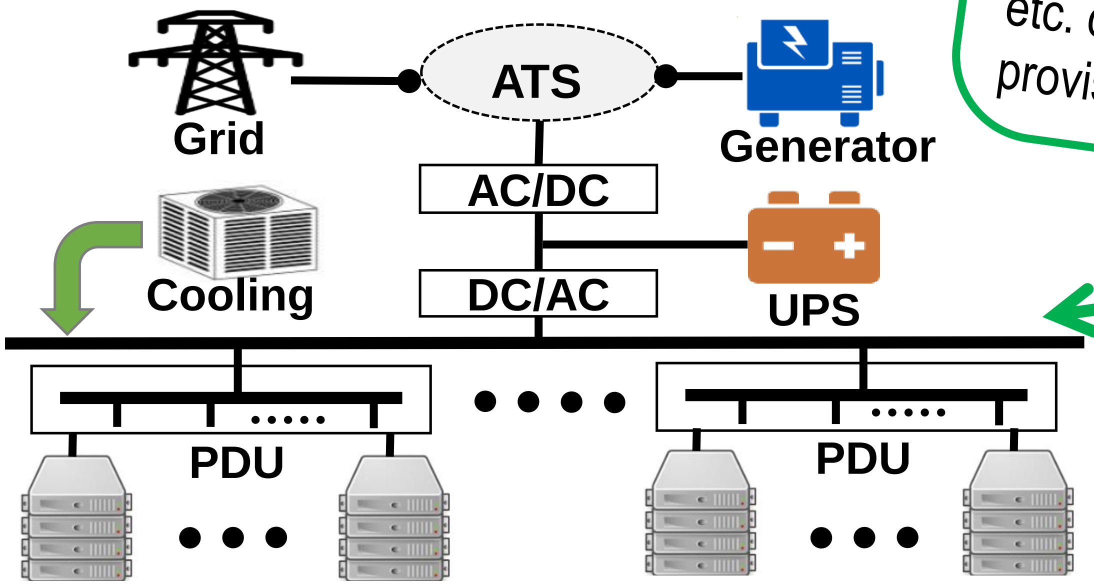

Diesel generator, UPS, PDU, cooling, etc. cost **$10-25 per Watt** of IT power provisioning

Infrastructure is really expensive
especially for **multi-tenant data centers**

Diesel generator, UPS, PDU, cooling, etc. cost **$10-25 per Watt** of IT power provisioning

Owned by operators

Owned by tenants

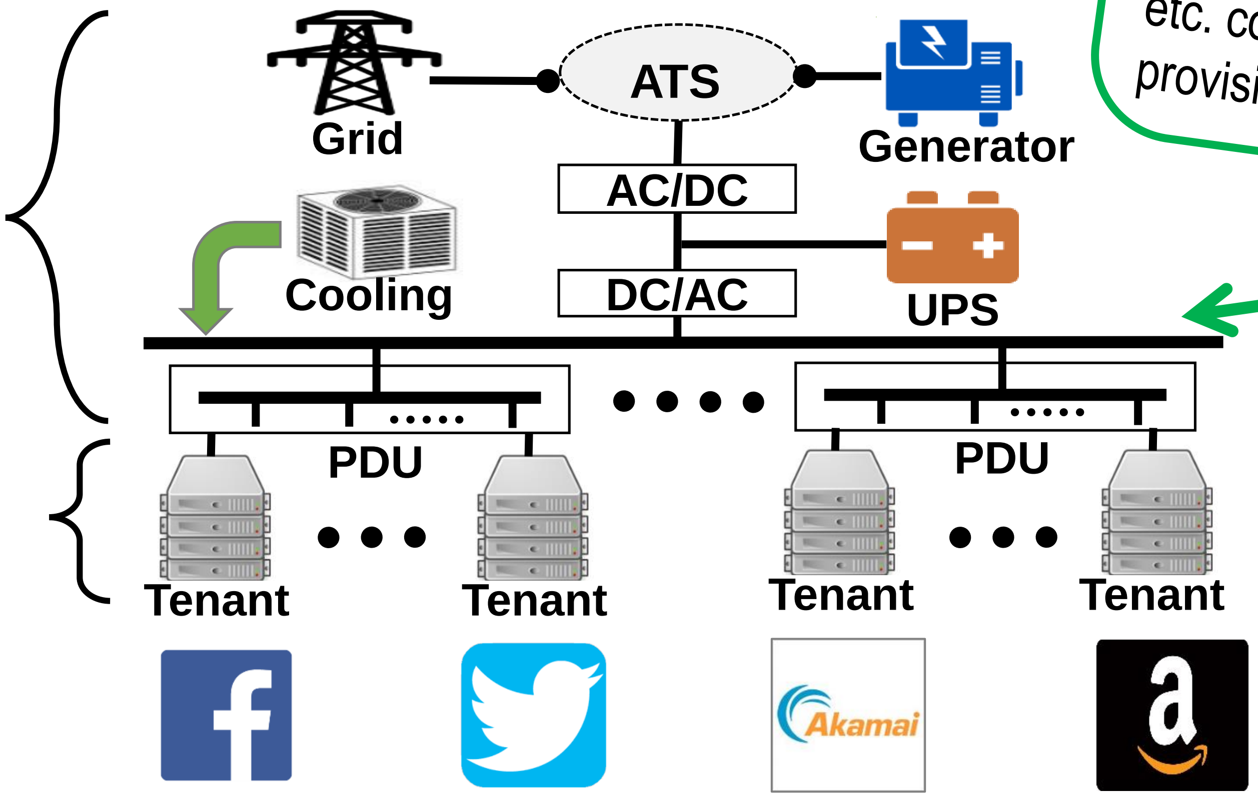

Infrastructure is really expensive especially for multi-tenant data centers

Owned by operators

Owned by tenants

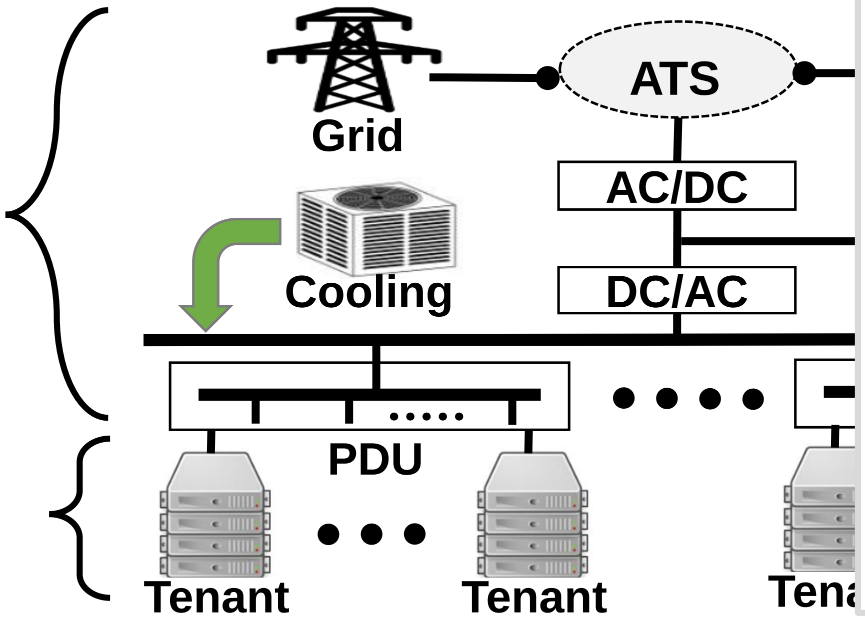

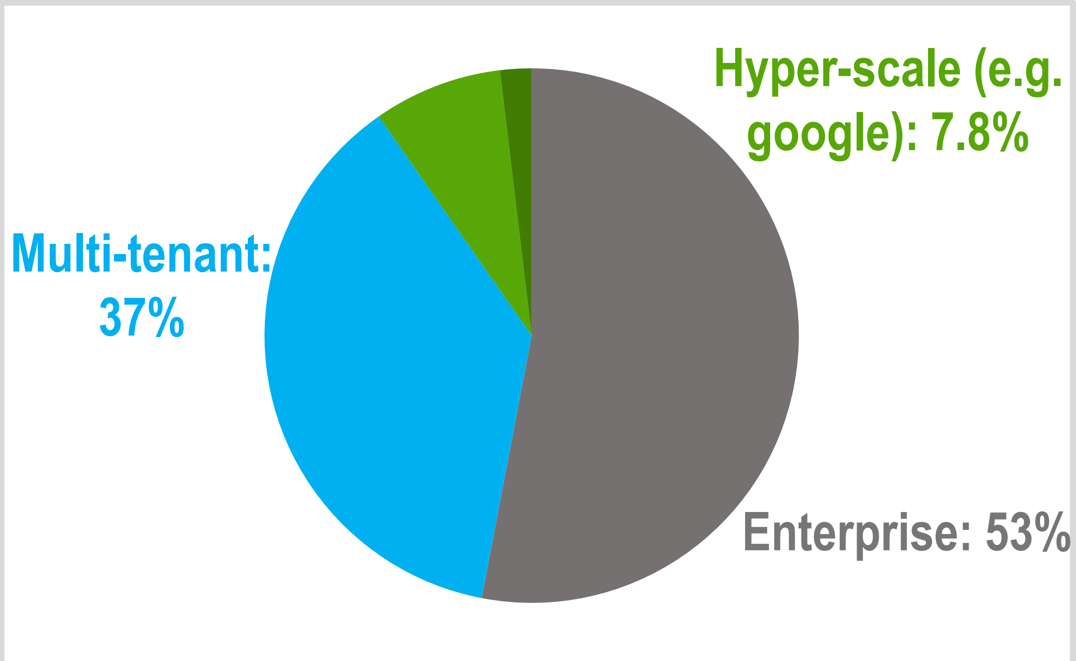

Percentage of total data center industry electricity usage

Pie Chart from CoreSite's "One Wilshire" (Photo: CoreSite)

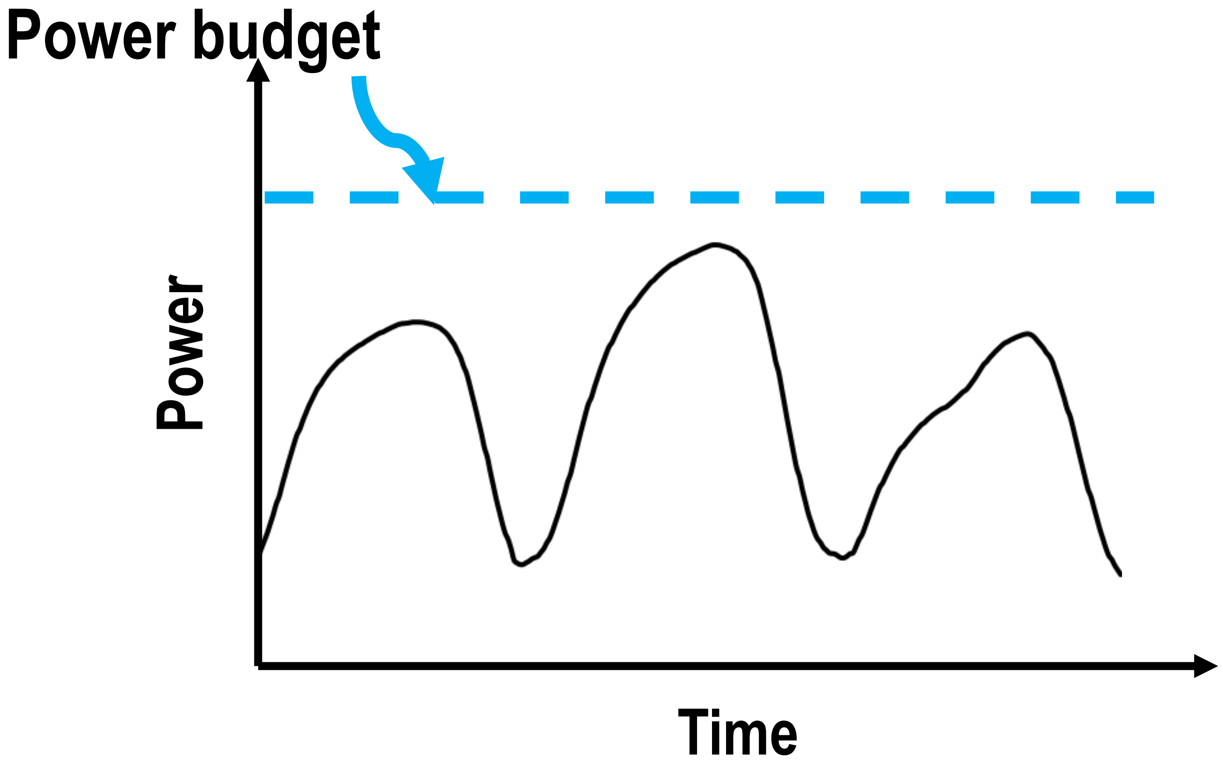
Power budget
Power
Time

We need to maximize the utilization!

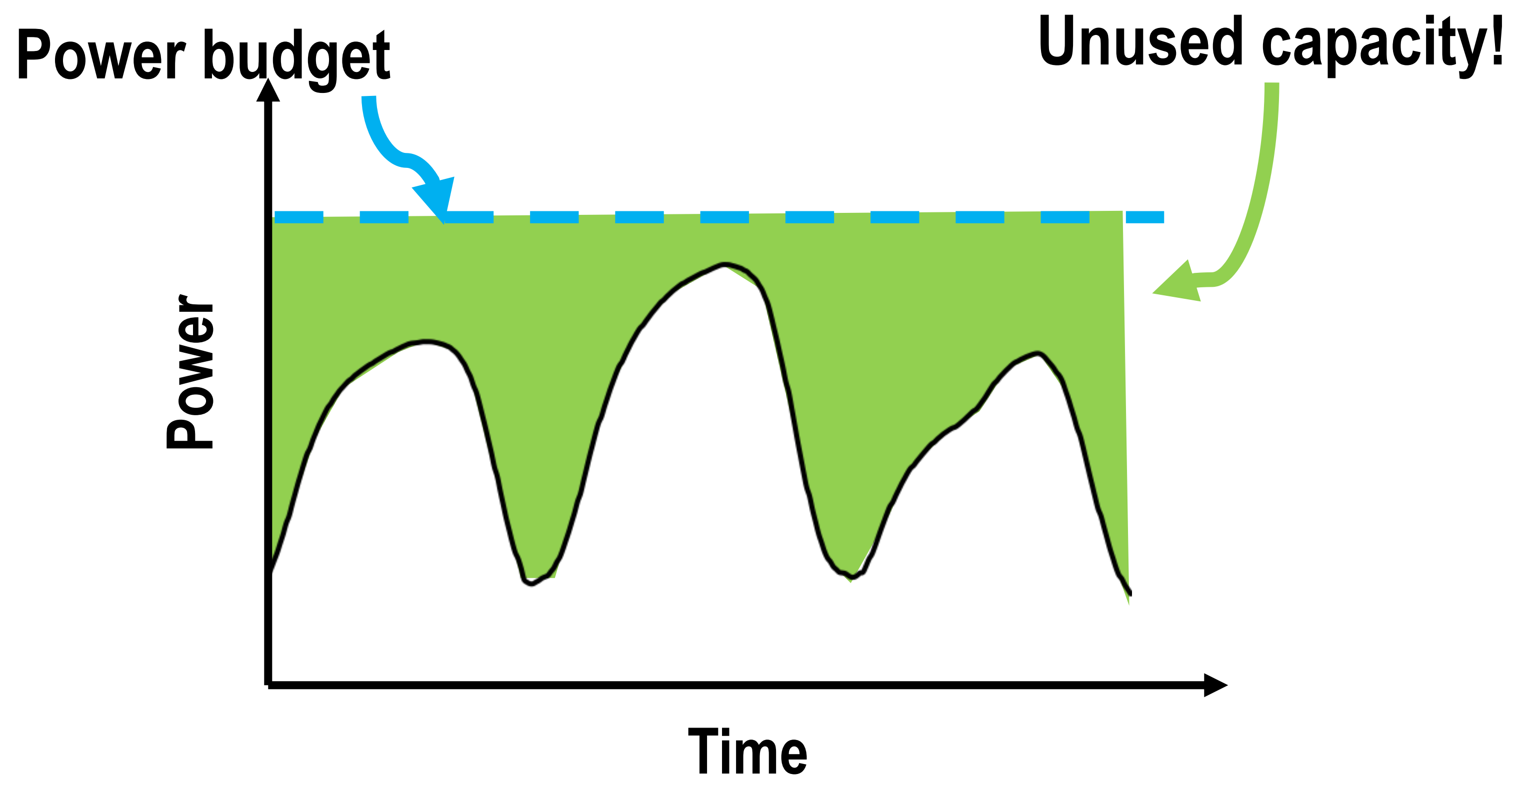

We need to maximize the utilization!

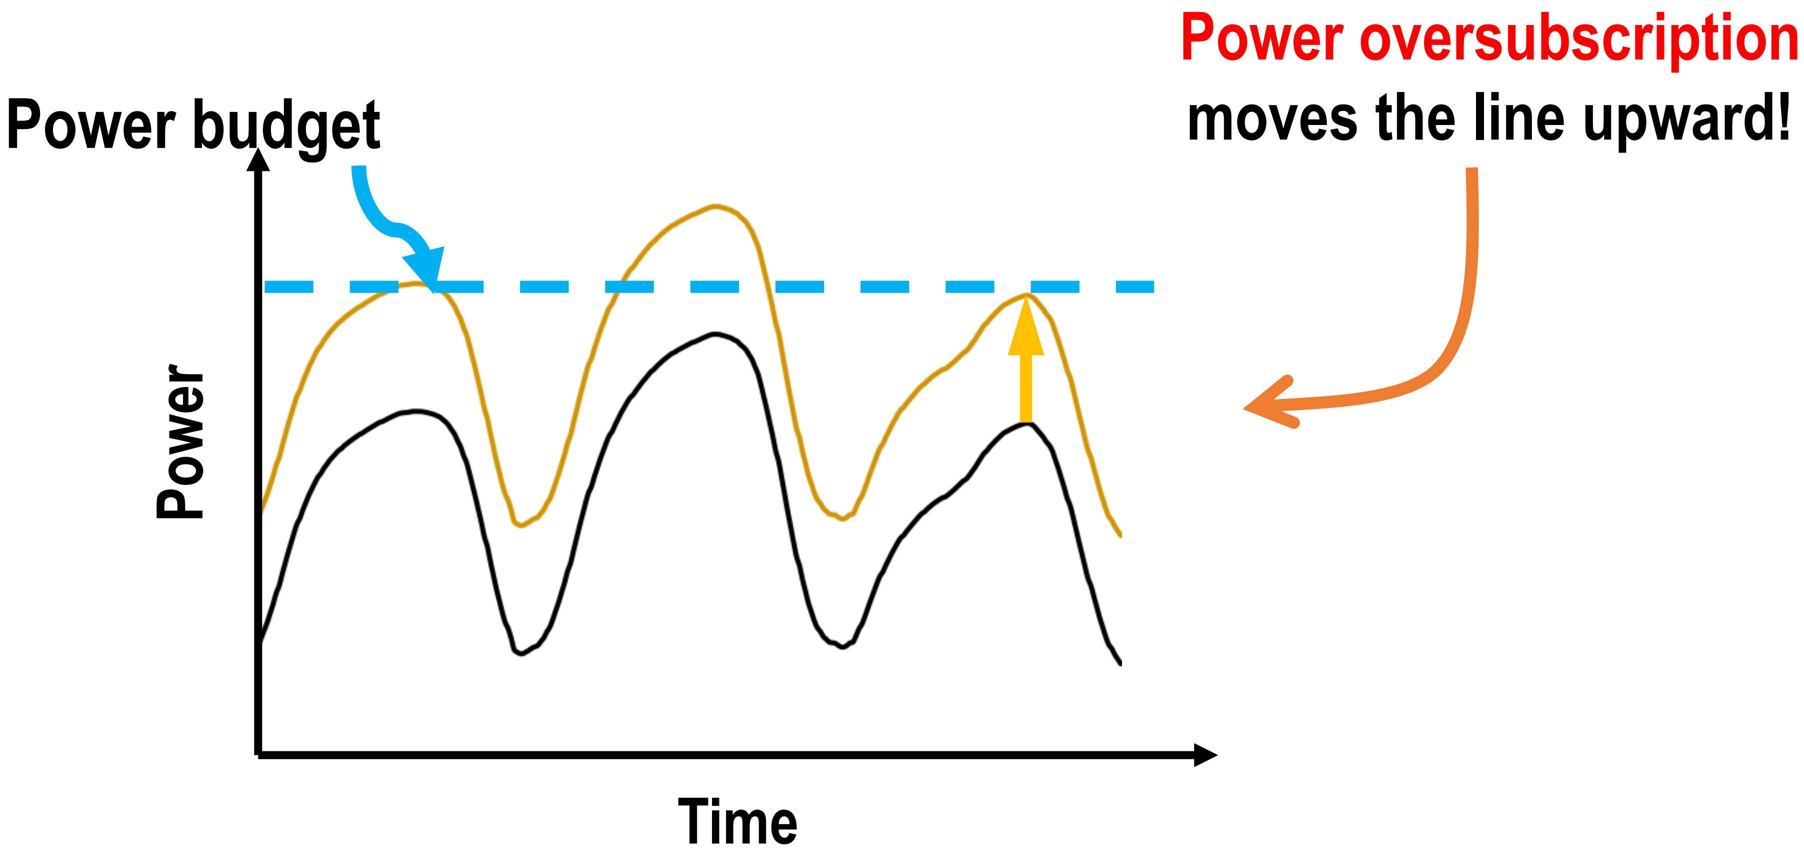

Benefits of power oversubscription

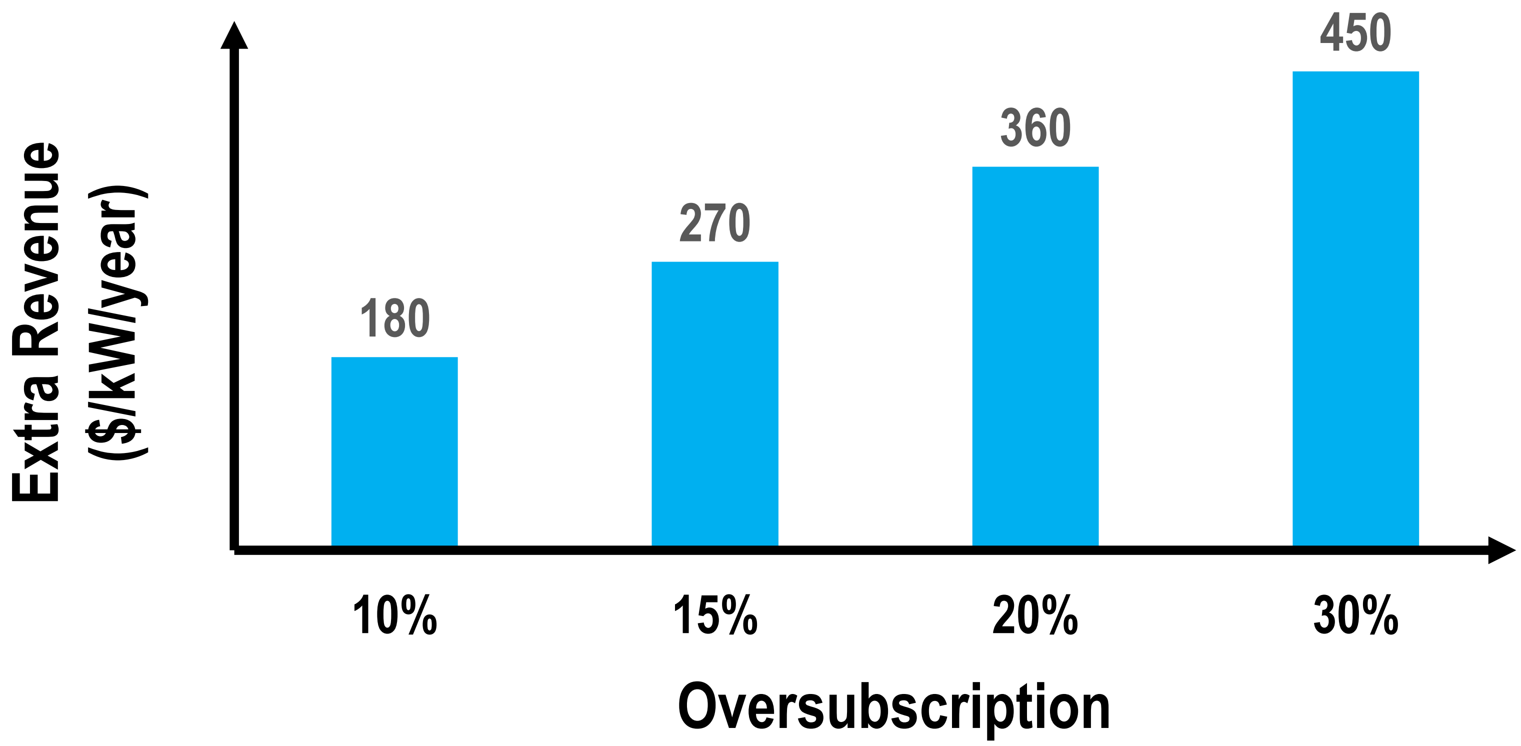

Challenges for power oversubscription

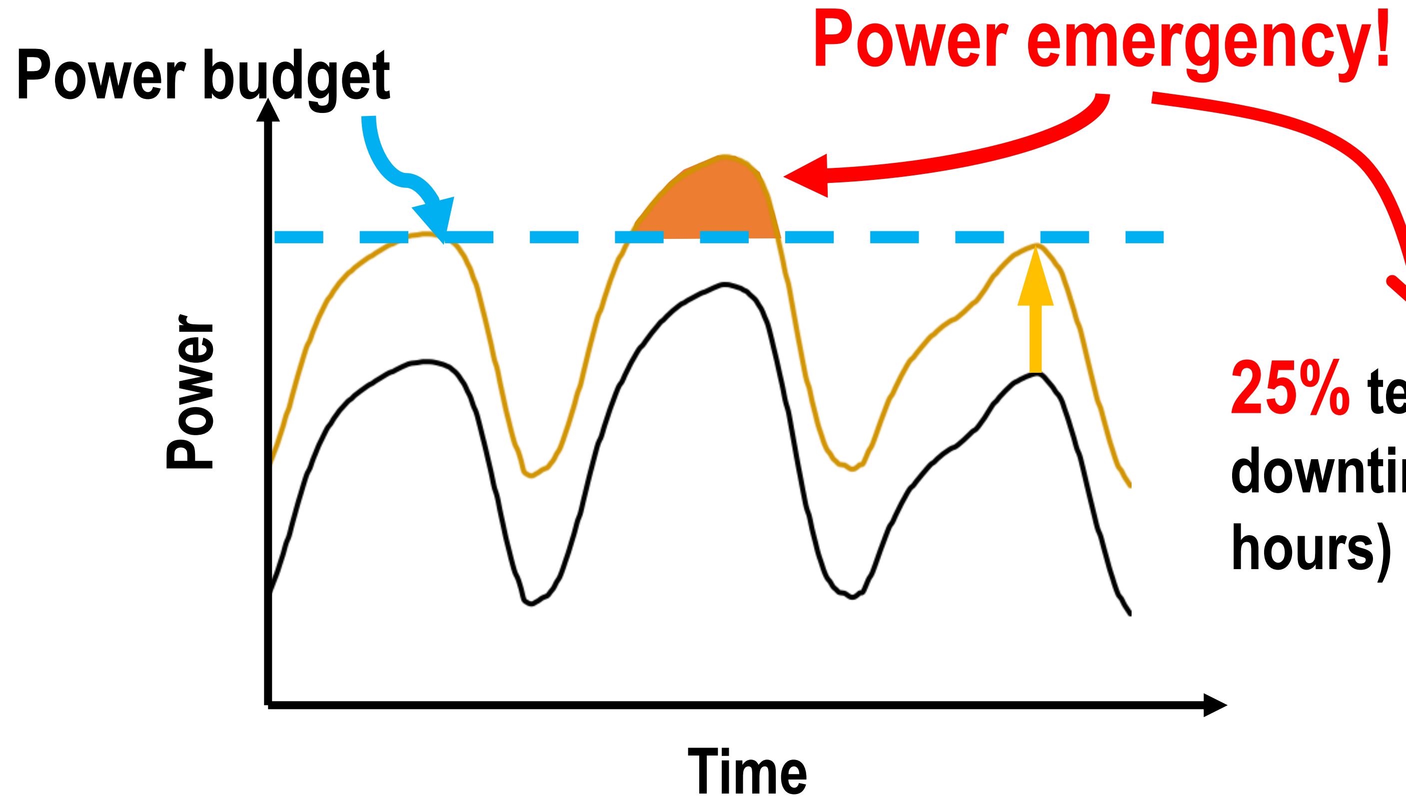

25% tenants ≥ 1 downtime (several hours) in 2014.

Challenges for power oversubscription

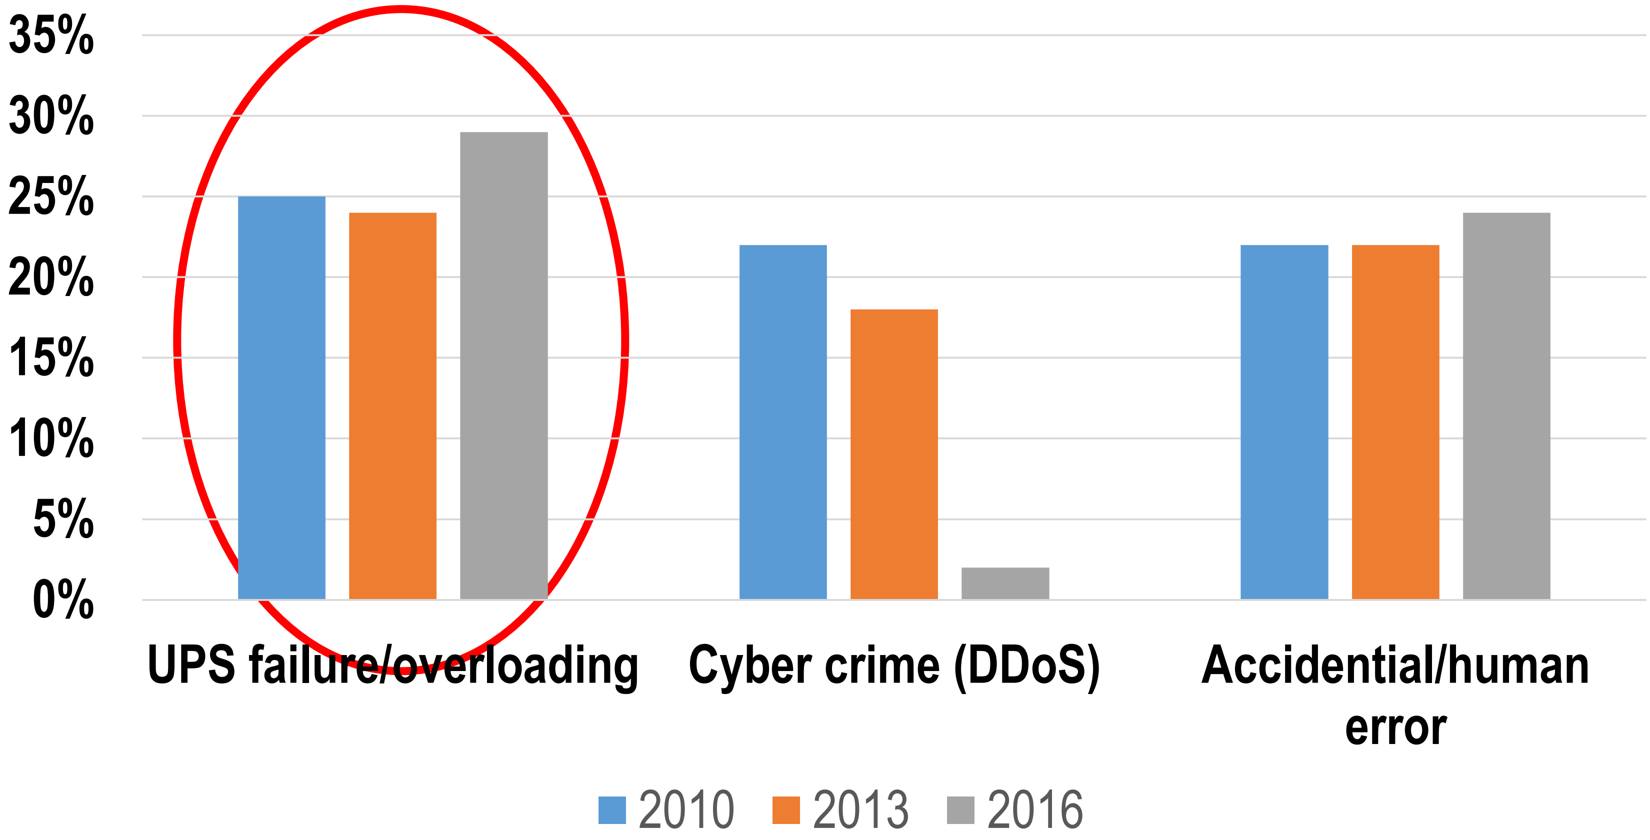

Data from report "Cost of Data Center Outages" by Ponemon Institute, Jan 2016.

How data center operators currently handle emergencies?

Before an outage occurs:

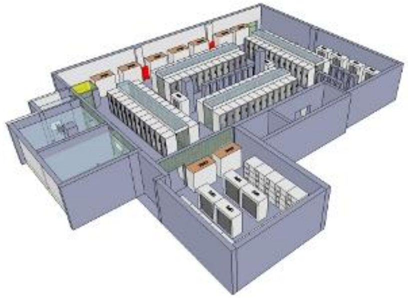

Operator

Nothing!

Tenants

How data center operators currently handle emergencies?

After an outage occurs:

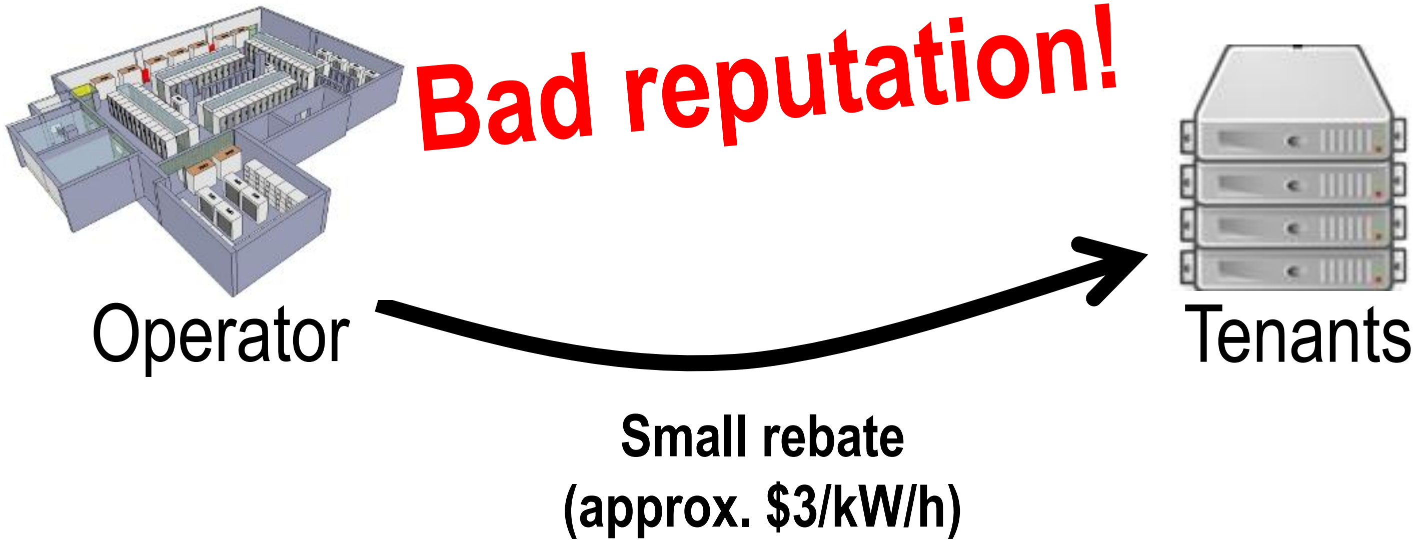

Consequences of power outage

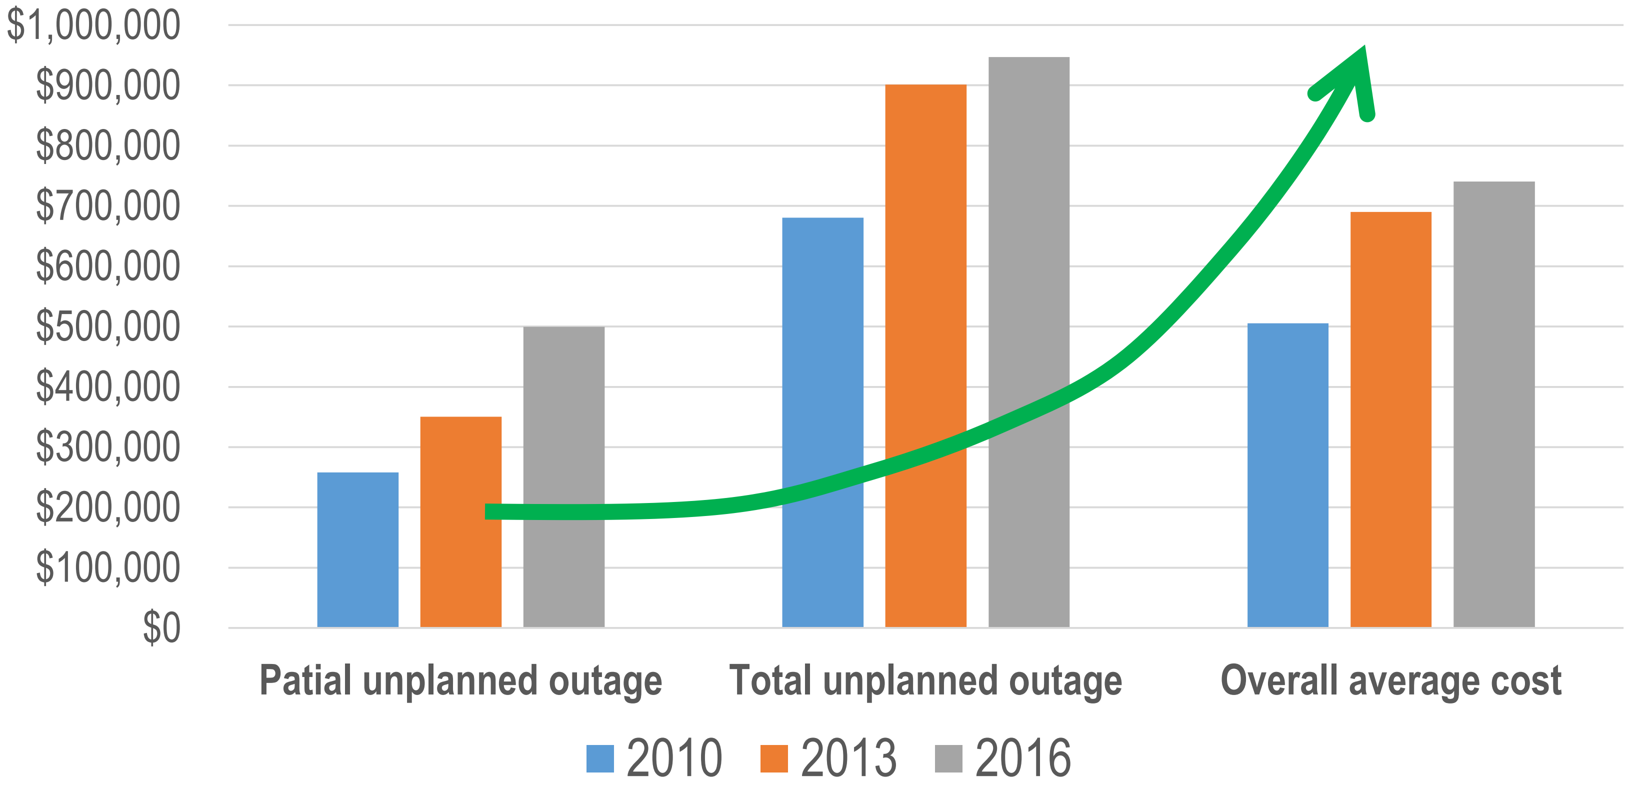

On average, each incident is a **million dollar loss**

Data from report “Cost of Data Center Outages” by Ponemon Institute, Jan 2016.

Consequences of power outage

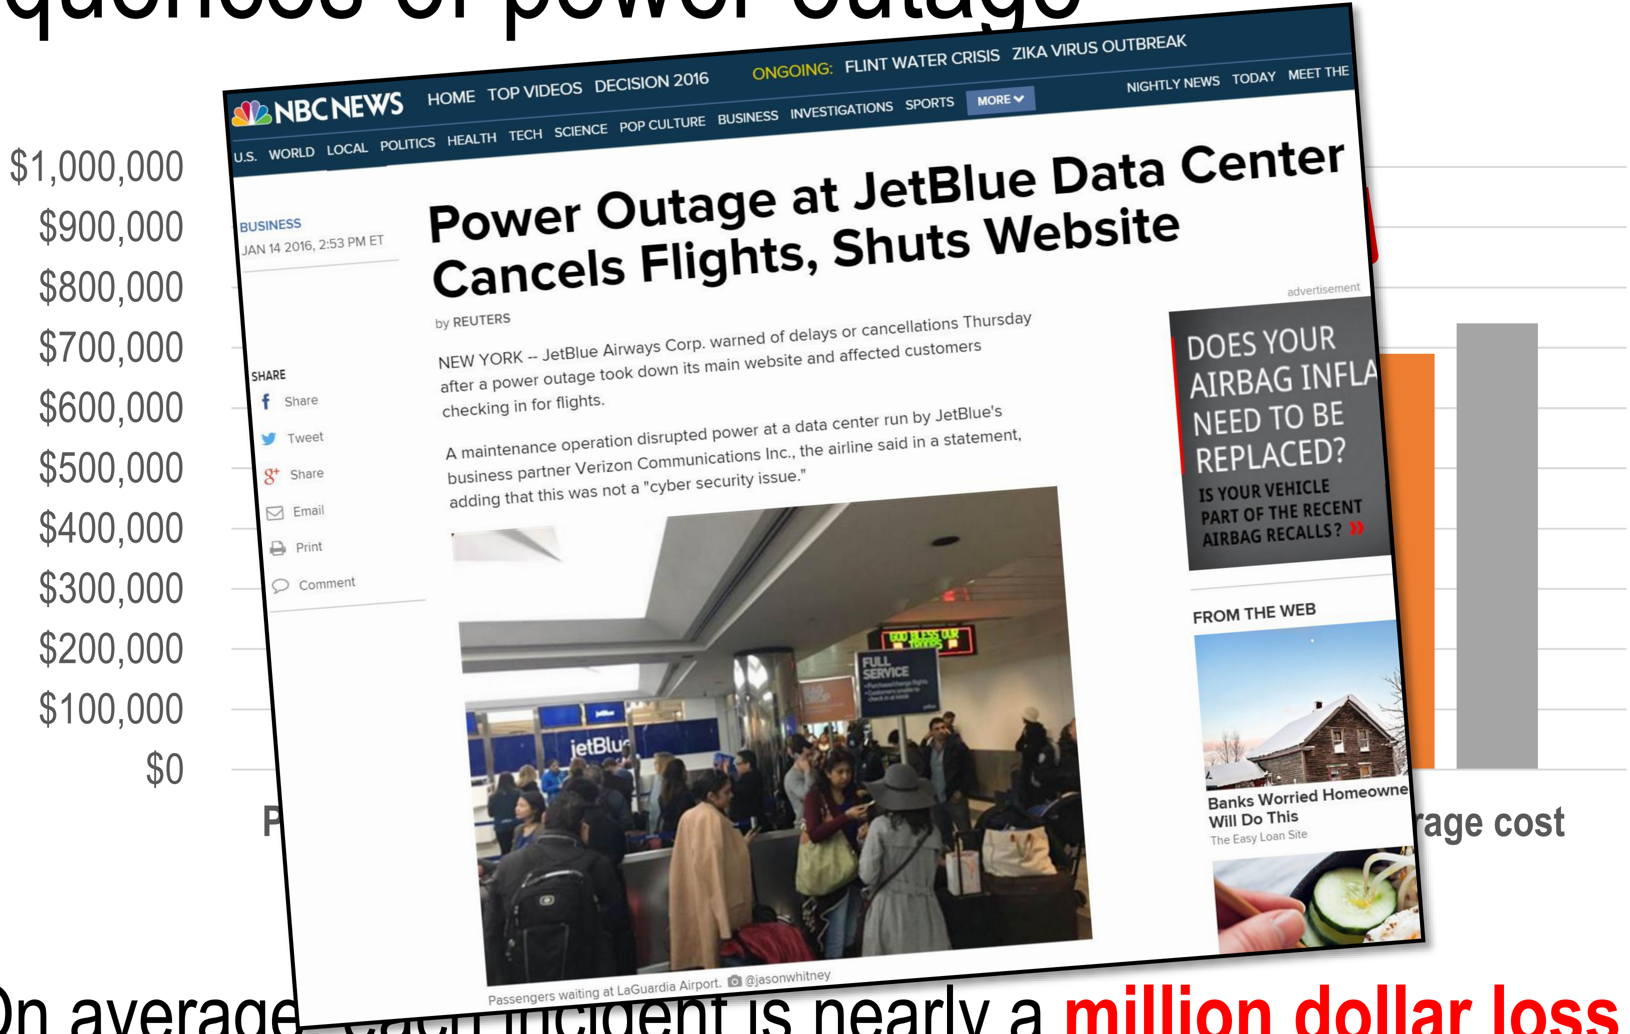

NBC NEWS HOME TOP VIDEOS DECISION 2016 ONGOING: FLINT WATER CRISIS ZIKA VIRUS OUTBREAK

NIGHTLY NEWS TODAY MEET THE

U.S. WORLD LOCAL POLITICS HEALTH TECH SCIENCE POP CULTURE BUSINESS INVESTIGATIONS SPORTS MORE

BUSINESS

JAN 14 2016, 2:53 PM ET

Power Outage at JetBlue Data Center Cancels Flights, Shuts Website

by REUTERS

SHARE

Share

Tweet

Share

Email

Print

Comment

NEW YORK -- JetBlue Airways Corp. warned of delays or cancellations Thursday after a power outage took down its main website and affected customers checking in for flights.

A maintenance operation disrupted power at a data center run by JetBlue's business partner Verizon Communications Inc., the airline said in a statement, adding that this was not a "cyber security issue."

Passengers waiting at LaGuardia Airport. @jasonwhitney

On average, each incident is nearly a **million dollar loss**

We need to handle power emergencies better!

Natural ideas

- Lower the IT power usage
  - There're many power capping solutions
    - DVFS, admission control, load migration, etc. [X. Wang, 2009][H. Lim 2011][X. Fu, 2011][A. Bhattacharya, 2012][D. Wang, 2013]
  - But, operator does NOT control tenants' servers
    - Even assuming it does, which tenants should reduce power and by how much?
- Static power reduction contracts
  - Cannot predict power reduction from tenants during an emergency

Natural ideas

- Lower the IT power usage
  - There’re many power capping solutions
    - … Fu,

Not applicable to multi-tenant data centers!

  - …
    - Even assuming it does, which tenants should reduce power and by how much?
- Static power reduction contracts
  - Cannot predict power reduction from tenants during an emergency

Goal: provide a **runtime** design to extract power reduction from **tenants** at minimum performance loss!

COOP: CO-Ordinated Power management

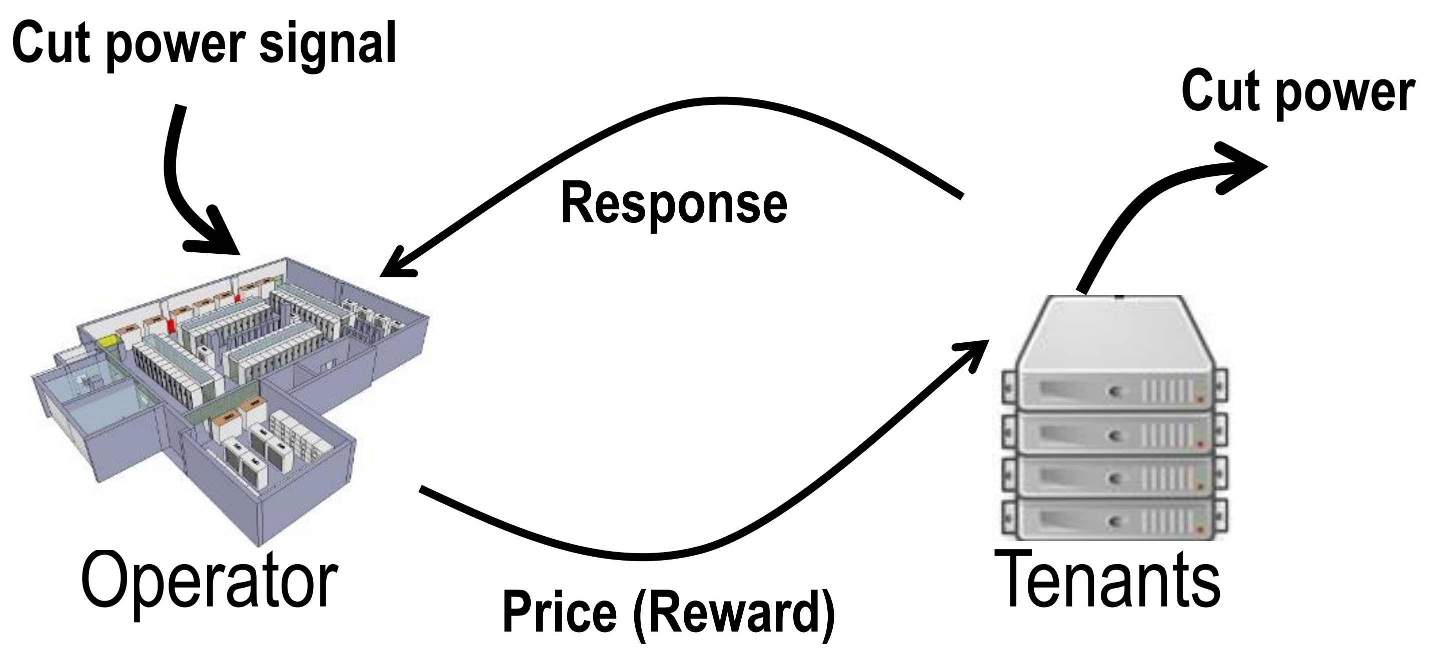

When a power emergency occurs…

- Two-level capping: high-level UPS and low-level PDU
  - UPS capacity exceeded by D_0
  - PDU capacity exceeded by D_i
  - N tenants: each cut power s_i and has a "performance cost" of $c_i(s_i)$

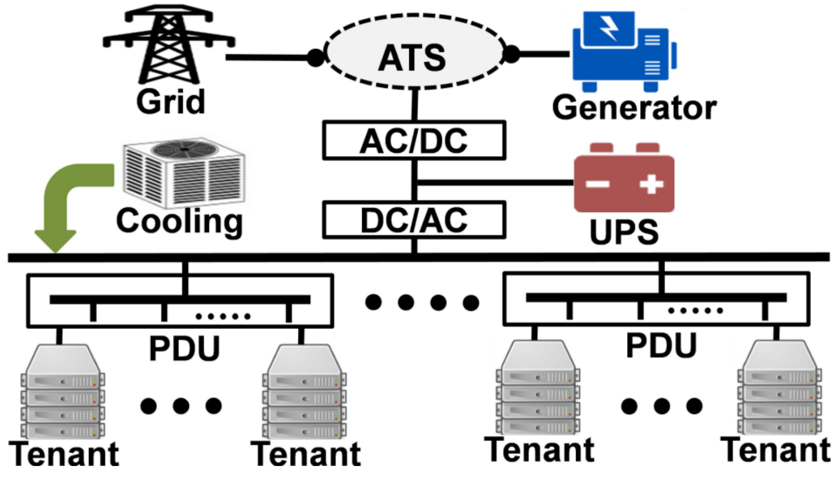

$$\min_{s_i \geq 0, i=1,2,\cdots,N} \sum_{i=1}^{N} c_i(s_i)$$

$$s.t., \quad \sum_{i \in \mathcal{N}_j} s_i \geq D_j, \text{ for } j = 0, 1, 2, \cdots, M,$$

How to solve it?

$$\min_{s_i \geq 0, i=1,2,\cdots,N} \sum_{i=1}^{N} c_i(s_i)$$

$$s.t., \quad \sum_{i \in \mathcal{N}_j} s_i \geq D_j, \text{ for } j = 0, 1, 2, \cdots, M,$$

- Centralized control doesn’t work…
- Market approach

Tenants report some, but not all, information via supply functions

Supply function bidding

$$\min_{s_i \geq 0, i=1,2,\cdots,N} \sum_{i=1}^{N} c_i(s_i)$$
$$s.t., \quad \sum_{i \in \mathcal{N}_j} s_i \geq D_j, \text{ for } j = 0,1,2,\cdots,M,$$

Pricing

Operator predicts tenants' responses; Tenants report **nothing** to the operator

Auction

Tenants report all information, i.e., "performance cost" $c_i(s_i)$; Operator sets prices accordingly.

Supply function $s(r)$

- If you offer me r, I will reduce power s_r…
  - Extensively studied in the context of electricity markets
- We choose a parameterized supply function as follows
  - Efficiency [R. Johari, 2011][N. Chen, 2015]

$$\boldsymbol{s_i(bi, r) = \left[\delta_i - \frac{b_i}{r}\right]^+}$$

Parameterized supply function bidding

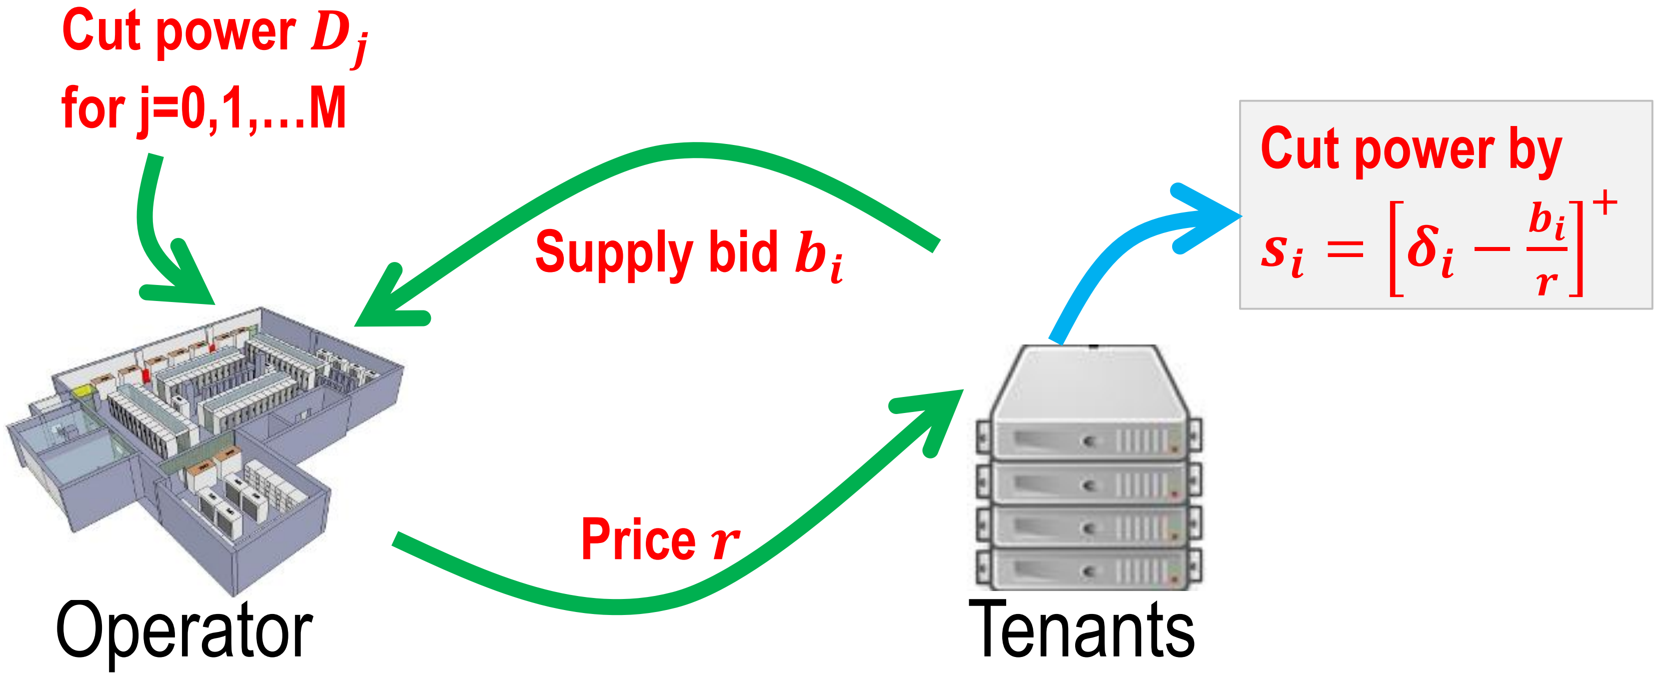

#1: Operator announces supply function $s_i(bi, r) = \left[\delta_i - \frac{b_i}{r}\right]^+$
#2: Tenant i submits bid b_i
#3: Operator clears market price r to satisfy multi-level power capping
#4: Power reduction is exercised

How to bid?

- Bid based on tenant's own performance cost, but **no need** to disclose it

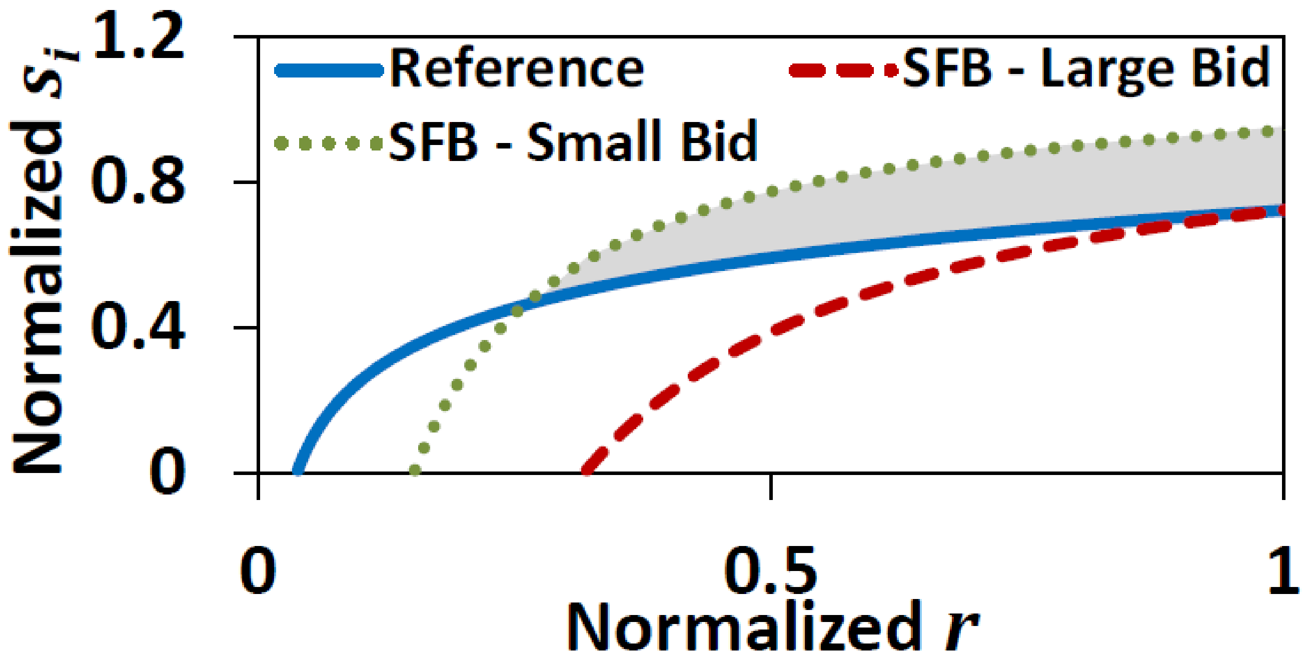

How to set price?

- Tenants reduce more power when offered higher price
- Just sufficiently large to make sure that tenants are reducing enough power
  - If no price is within the expected range (to ensure no profit loss for operator), then enter "failover" mode

$$\min \quad r$$

$$\text{s.t.} \sum_{i \in \mathcal{N}_{|}} s_i \geq D_j, \text{for } j = 0, 1, 2, \ldots, M$$

Implementation

Algorithm 1 COOP: Coordinated Power Management

1: Input: UPS and PDU capacities P_i^{cap} for $i = 0, 1, \cdots, M$
2: Monitor UPS and PDU power $P_i(t)$ continuously.
3: **if** $P_i(t) > P_i^{cap}$ for any $i = 0, 1, \cdots, M$ **then**
4:     Start waiting timer T_w
5: **end if**
6: **while** T_w has not expired **do**
7:     **if** $P_i(t) \leq P_i^{cap}$ for all $i = 0, 1, \cdots, M$ **then**
8:         Go back to Line 2
9:     **end if**
10: **end while**
11: ▷ *Entering "power capping" mode*
12: **if** $P_i(t) > P_i^{cap}$ for any $i = 0, 1, \cdots, M$ **then**
13:     Set $D_i \leftarrow [P_i(t) - P_i^{cap}]^+$
14:     Announce $s_i(b_i, r) = [\delta_i - \frac{b_i}{r}]^+$ to tenant i
15:     Tenant i decides its bid b_i
16:     Set price $r = \min_{r'} \{r' \in [r_{\min}, r_{\max}] \mid \sum_{i \in \mathcal{N}_j} s_i(b_i, r')$
17:     $\geq D_j$, for $j = 0, 1, \cdots, M\}$
18:     Each tenant i reduces $s_i(b_i, r)$ power
19: **end if**
20: ▷ *Leaving "power capping" mode*
21: **wait** until $P_i(t) \leq P_i^{cap} - D_i$ for all $i = 0, 1, \cdots, M$
22: Start capping timer T_c and **wait** until T_c expires or $P_i(t) > P_i^{cap} - D_i$ for any $i = 0, 1, \cdots, M$
23: **if** $P_i(t) > P_i^{cap} - D_i$ for any $i = 0, 1, \cdots, M$ **then**
24:     Go back to Line 21
25: **end if**
26: **if** T_c expires **then**
27:     Notify tenants to resume normal operation
28:     Calculate the power capping duration T_o
29:     Provide tenant i with a reward of $z_i = T_o \cdot r \cdot s_i$
30:     Go back to Line 2
31: **end if**

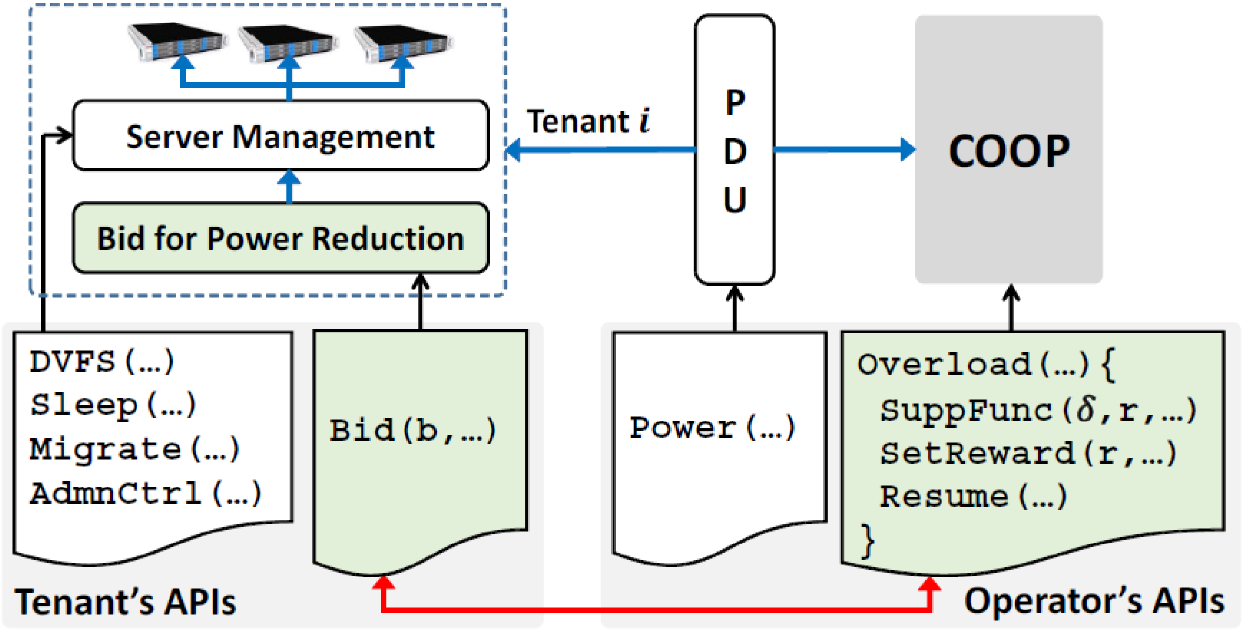

Evaluation Methodology

Tenant	Type	No. of Servers	Tenant's Max. Power	Location	Cluster's Max. Power
#1	Web search	2	200 W	Cluster#A	740 W
#2	KVS	2	310 W		
#3	Hadoop	2	230 W		
#4	Web search	3	300 W	Cluster#B	530 W
#5	Hadoop	2	230 W		

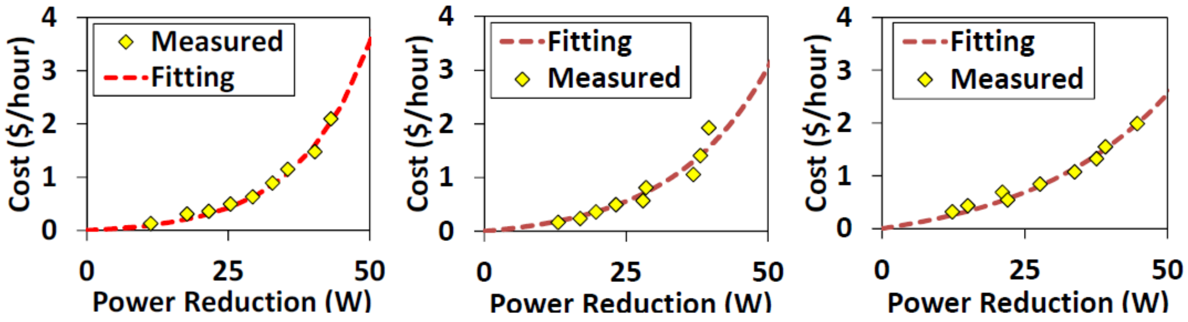

(a) T#1- Web search (b) T#2- KVS (c) T#3- Hadoop

- 5 tenants running different workloads housed on two clusters
- DVFS for power reduction

COOP is close to Optimal

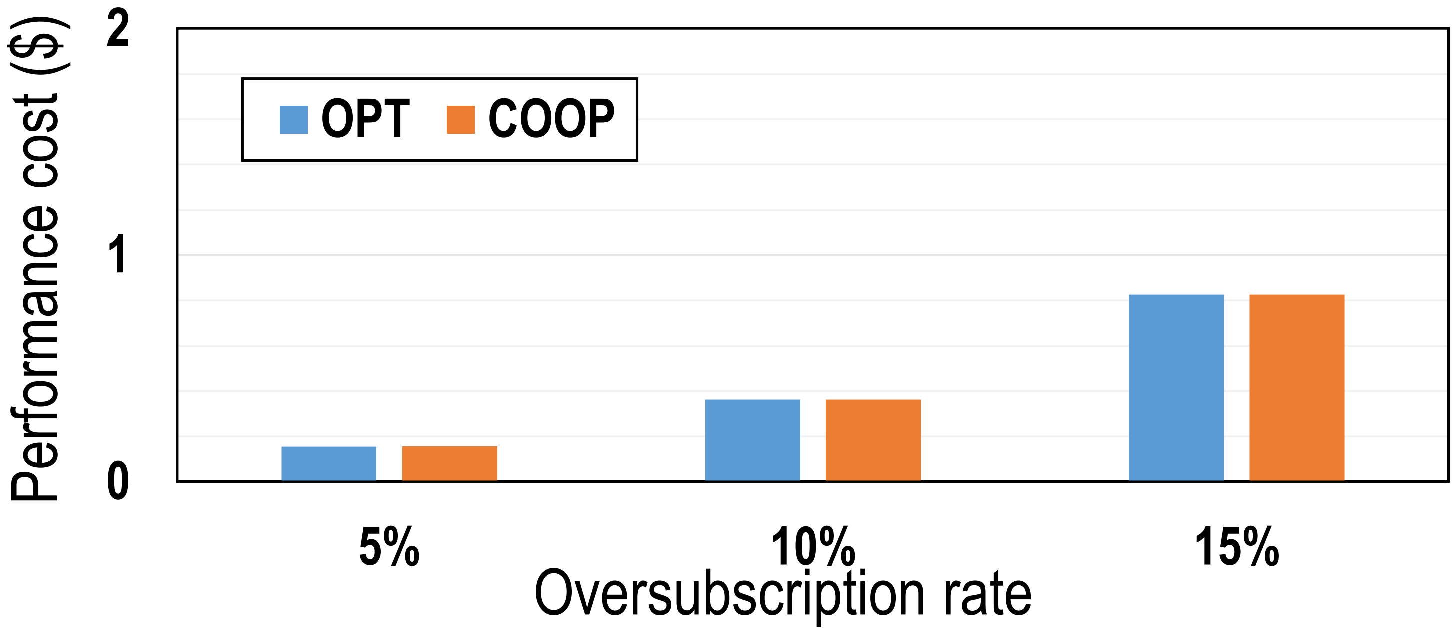

- COOP almost minimizes the performance costs as OPT
  - OPT is an idealized case where the operator dictates tenants' power reduction as in an owner-operated data center
  - Settling time: just <1 second

COOP is win-win

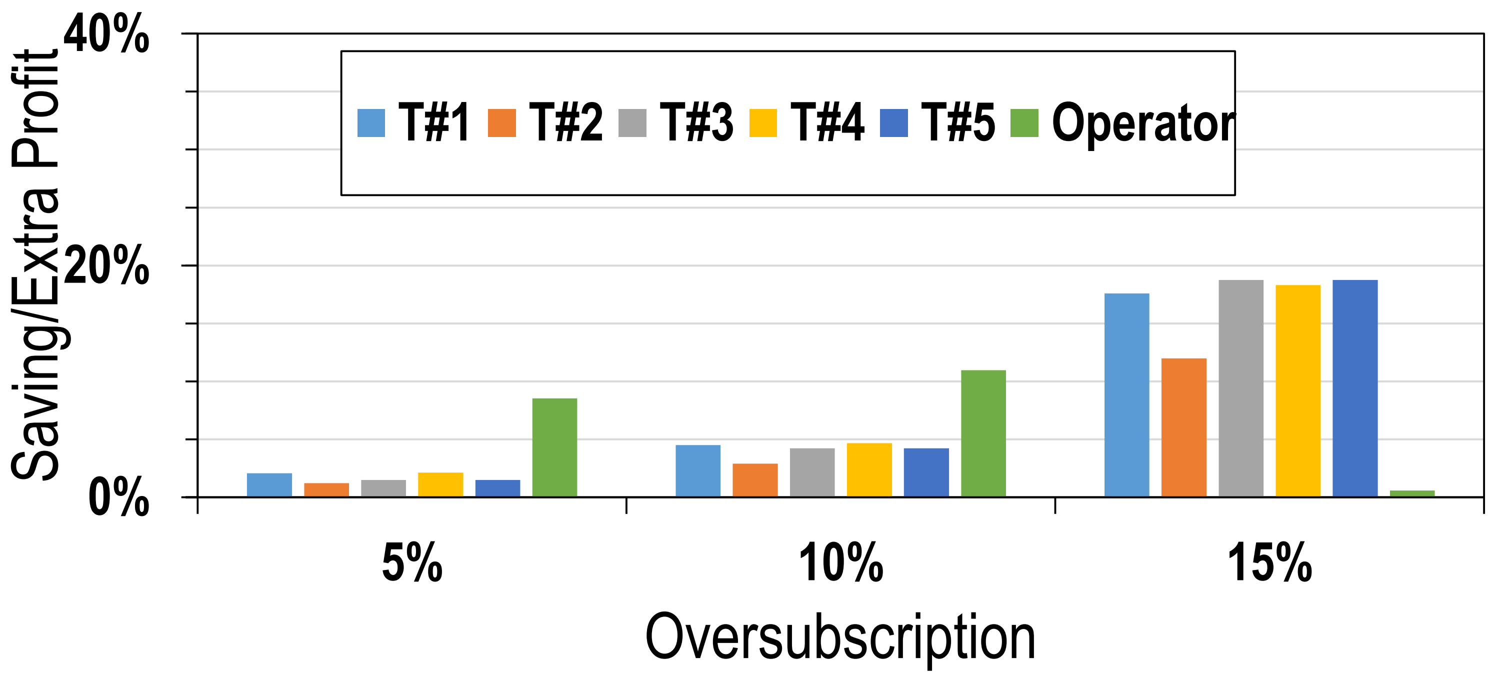

- Tenants reduce power cost with minimum (temporary) performance impact
- Operator increases profit by selling capacity to more tenants

COOP: CO-Ordinated Power management

A **market-based** approach for handling power emergencies and helping operator better oversubscribe data center capacity

Simple, Scalable & Efficient