2016 ACM/IEEE 43rd Annual International Symposium on Computer Architecture

# Power Attack Defense: Securing Battery-Backed Data Centers

Chao Li[1] Zhenhua Wang[2] Xiaofeng Hou[1] Haopeng Chen[2] Xiaoyao Liang[1] Minyi Guo[1]

[1]Department of Computer Science and Engineering, Shanghai Jiao Tong University, Shanghai, China
[2]School of Software, Shanghai Jiao Tong University, Shanghai, China
{lichao, liang-xy, guo-my} @cs.sjtu.edu.cn

***Abstract*** — ***Battery systems are crucial components for mission-critical data centers. Without secure energy backup, existing under-provisioned data centers are largely unguarded targets for cyber criminals. Particularly for today's scale-out servers, power oversubscription unavoidably taxes a data center's backup energy resources, leaving very little room for dealing with emergency. Besides, the emerging trend towards deploying distributed energy storage architecture causes the associated energy backup of each rack to shrink, making servers vulnerable to power anomalies. As a result, an attacker can generate power peaks to easily crash or disrupt a power-constrained system. This study aims at securing data centers from malicious loads that seek to drain their precious energy storage and overload server racks without prior detection. We term such load as Power Virus (PV) and demonstrate its basic two-phase attacking model and characterize its behaviors on real systems. The PV can learn the victim rack's battery characteristics by disguising as benign loads. Once gaining enough information, the PV can be mutated to generate hidden power spikes that have a high chance to overload the system. To defend against PV, we propose power attack defense (PAD), a novel energy management patch built on lightweight software and hardware mechanisms. PAD not only increases the attacking cost considerably by hiding vulnerable racks from visible spikes, it also strengthens the last line of defense against hidden spikes. Using Google cluster traces we show that PAD can effectively raise the bar of a successful power attack: compared to prior arts, it increases the data center survival time by 1.6~11X and provides better performance guarantee. It enables modern data centers to safely exploit the benefits that power oversubscription may provide, with the slightest cost overhead.***

***Keywords- data center; battery; power attack; defense;***

## I. INTRODUCTION

Data center servers are becoming tightly coupled with, and more dependent on, local energy storage devices. In recent years we have witnessed a considerable interest in deploying massive distributed energy backup (DEB). For example, Google and Facebook have started to explore small-scale battery backup units in each rack or chassis to reduce power conversion loss and facility footprint [1, 2]. According to Microsoft, its newly released distributed local energy storage (LSE) will bring up to 15% improvement in power usage effectiveness (PUE) and up to 5X cost reduction over a central, bulky UPS system [3]. At Hitachi, researchers have demonstrated that in-rack DEB design could improve the already impressive efficiency of an intelligent data center power distribution system by over 8% [4]. Today, per-server battery backup unit is available from many vendors such as HP and Quanta [5, 6].

Looking ahead, distributed battery holds great promise in high-performance data center design. It is not only a more energy-efficient alternative to current uninterruptible power supply (UPS) system, but also easy to scale and maintain [7, 8]. It could eliminate a potential single point of failure (SPOF) that centralized UPS systems may have [7]. More importantly, a DEB-based data center is able to oversubscribe the power infrastructure without affecting server performance. The occasional power demand peaks can be shaved by a fraction of battery units effectively and no performance capping is performed [7-9].

Despite the above advantages, the power and energy related security issue has become the Achilles' heel of a DEB-based data center. Without obtaining a privileged access, an attacker can gain key energy backup information through various side-channels. Given the growing flexibility of Internet service and potential bugs of cloud APIs, a malicious load can abuse the power and energy resources (especially stored backup energy) in a data center [10]. For example, by creating excessive floating-point operations or triggering more cache misses, the attacker can increase system resource consumption considerably. The malicious load, which we refer to as *Power Virus*, is able to generate simultaneously occurred power surge to overload the system [10, 11]. Specifically, a power virus can first create *non-offending visible power peaks* (disguised as benign loads) to drain energy backup and then be mutated to create *offending hidden power spikes* that can bring down the victim rack without prior detection.

As we transition from centralized to distributed energy storage architecture, server racks unfortunately become vulnerable to power virus. Local power failure is more prone to occur since DEB units physically lack the capacity for handling extended outage duration (e.g., less than 2 minutes under full load [2, 6]). Worse, the DEB architecture often presents a ready-made "divide and conquer" solution for attackers — creating a local power peak is much easier than overloading the entire data center.

The security issue turns out to be particularly acute in many data centers that are heavily power-constrained. To save the significant power cost, it is not unusual that data centers aggressively oversubscribe their power systems [12, 13]. As companies continue to squeeze more servers into their existing data center, the risk of power violation is rising

1063-6897/16 $31.00 © 2016 IEEE
DOI 10.1109/ISCA.2016.50

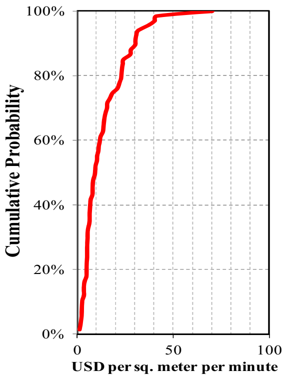

**Fig. 1.** The cumulative distribution function (CDF) of datacenter power failure cost [19]. The cost covers detection and recovery

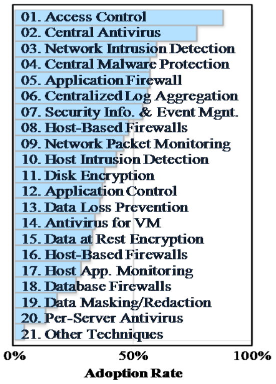

**Fig. 2.** Survey of data center security technologies [20]. Energy/power related security issues demand more attention

rapidly. In addition, DEBs have been frequently used as energy buffer in recent green data center designs to handle the power variability [14-17]. In both cases, batteries often experience unusual cyclic usage but do not receive timely recharge. Without enough backup energy, racks are left unguarded from malicious loads. By far the biggest root cause of power outage is battery failure and capacity exceeded [18], which could have been avoided with a proactive security-aware energy management.

Oftentimes, the power-related attack could have devastating effects on the victim data centers. It can cause service interruption on the blackout servers and even irreparable financial loss to an organization. Unplanned power outage has been shown to cost over $10 per square meter per minute for 40% of the benchmarked data centers (Figure 1) [19]. On average, the financial loss of a data center power outage in 2013 is more than $7900 per minute — an increase of 40% compared to 2010 [19]. According to a recent survey, more than 75% data centers require at least 2 hours to investigate and remediate incidents [20]. It means that a successful power-related attack can easily cause the victim data center to lose one million dollars.

While various technologies are available to protect our servers, the security issue associated with energy/power has been largely overlooked today (Figure 2). It is very hard to defend against power-related attacks indirectly with existing methods such as infusion detection and access control. This is because power analysis based on load statistics is often resource-consuming and the results are often inaccurate [11]. In fact, over 70% data center operators in a large-scale survey believe that their monitoring programs lack the fine-grained visibility at the server level [20]. Although advanced power metering can be used to allow for real-time analysis, it is not available in most data centers. Fine-grained sampling and metering are also prohibitive since it requires costly implementation of per-server metering. As a result, attackers can often manage to launch power virus without prior detection [6].

This study aims to understand the vulnerability of battery-dependent data centers and provides an initial solution. Specifically, we focus on the question of *how massive data center battery units can be gracefully tamed and leveraged to tackle the difficult and costly challenge posed by malicious loads*. This is crucial to companies who want to exploit DEB to improve energy-/cost- efficiency but cannot afford to compromise service availability.

We propose *power attack defense* (PAD), a novel design patch for securing data centers backed by distributed battery units. The salient feature of PAD is two-fold: 1) it does not require a detailed knowledge of the underlying workloads; 2) it is built on lightweight software and hardware mechanisms. Specifically, PAD provides an additional layer of safety in data centers through a novel two-phase power diagnosis and management.

In the first phase, PAD handles the visible peaks through software scheduling. Rather than treats each DEB as separate energy backup, PAD creates a virtual battery pool called *vDEB* to enable load sharing among spatially dispersed battery units. It leverages the power budget enforcing capability of today's intelligent PDUs to adjust DEB utilization of each rack. This proactive maintenance keeps massive DEB units operating in a coordinated manner, thereby avoiding vulnerable servers.

In the second phase, PAD uses a multi-layer DEB architecture to handle the more dangerous hidden spikes. In contrast to prior work that only have homogeneous peak shaving DEB units, our approach leverages a dedicated small-scale battery called *μDEB* at the rack-level to assist the server-level DEBs. It can automatically shave power spikes to avoid circuit breaker tripping.

PAD slightly increases software/hardware complexity but brings attractive security benefits. It does not require significant modifications on state-of-the-art designs: *vDEB* can be implemented as a small service extension to existing battery management mechanisms and *μDEB* is largely built upon existing energy storage components to provide better security guarantee. Our evaluation shows that PAD has minor performance/cost overhead (Section 6).

In this paper we introduce and motivate security-based design for power/energy management in data centers. We build a prototype to demonstrate the vulnerability of untamed batteries and we simulate different attacking scenarios at a larger scale using a Google compute cluster trace collected from over 220 machines in one month [21].

To the best of our knowledge, this work reflects the first systematic study of the security issue (with an emphasis on energy/power) in emerging battery-backed data centers.

This paper makes the following contributions:

- We describe a general threat model for power-related attacks (i.e. the creation of *power virus*). We discuss how a sophisticated attacker can leverage *visible peaks* and *hidden spikes* to bring down a DEB-based data center and characterize representative attacking scenarios in detail.

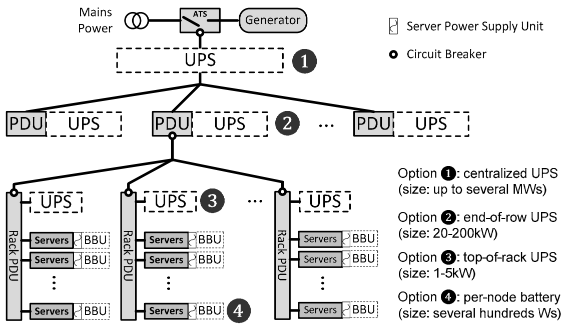

**Fig. 3.** Different battery deployment methods in a data center

- We present a three-level security policy for future DEB-based mission-critical data centers. We propose *power attack defense* (PAD), a novel light-weight software and hardware design patch for shielding vulnerable data center server racks from potential power-related attacks.
- We thoroughly evaluate the effectiveness of PAD through both real-system measurement and Google trace simulation. We show that PAD can effectively improve the critical data center survival time by 1.6-11X. It can provide better performance guarantee and has minor cost overhead.

The rest of this paper is organized as follows. §2 introduces background. §3 demonstrates our threat model. §4 proposes PAD. §5 describes experimental methodology. §6 presents evaluation results. Finally, §7 discusses related work and §8 concludes this paper.

## II. Background

Modern data centers normally include complex power provisioning systems that have very stringent capacity constraints. Safely oversubscribing the power infrastructure has become a critical need in data centers today due to the very high power capacity cost and power outage cost.

### A. Battery Backup Infrastructure

Currently there are primarily four ways to deploy batteries in a data center (Figure 3). To avoid the very costly over-provision of battery capacity, normally only one of the four backup methods is used. The size of each battery unit varies from hundreds watts to several MWs.

Conventional data centers mainly rely on bulk UPS (uninterruptible power supply) battery to provide the interim backup power in case of utility power outages. The UPS typically locates between the data center-level power panels and the cluster-level power distribution units (PDU). In some cases, PDU-level UPSs may be deployed to improve reliability. Most online UPSs need to convert power twice: input (AC to DC) and output (DC to AC) conversion. This can result in significant power loss.

Distributed energy backup (DEB) devices arise due to the needs to improve data center efficiency and facilitate peak power shaving [7]. They can be installed as top-of-rack UPS (or a battery cabinet next to the rack) or in each server enclose. Because they are DC based, the energy efficiency can be greatly improved. By directly integrating battery units locally and using DC voltage as backup, one can eliminate double-conversion. More importantly, one can easily switch a fraction of server racks to their local energy storage to shave/hide the power peaks at the data center level. A central UPS system cannot be used to support a fraction of data center servers: it either takes over the entire data center or serves as an idle power backup.

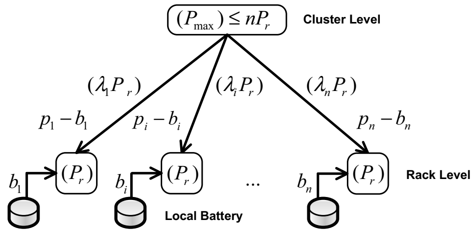

**Fig. 4.** Power oversubscription model. The maximum power budget of each node/edge is given in parentheses

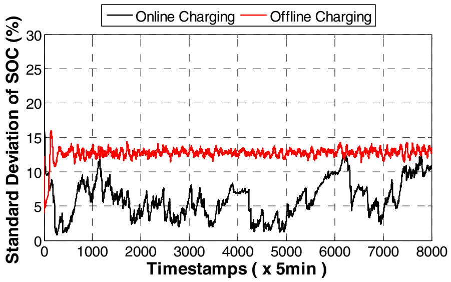

**Fig. 5.** Uneven utilization of distributed battery system

### B. Power Oversubscription Model

The power infrastructure is often one of the most expensive and longest lead time items in data center design, ranging between $10~25 per watt [22]. In addition, utility companies charge an additional power demand fee at many dollars per kW [23] every billing cycle. As a result, data centers normally over-provision their server systems in order to achieve the best return on investment (ROI).

In this study we focus on a typical two-stage power distribution at the server cluster level (Figure 4). Assume that the power demand of each rack is $p_i$, local batteries are responsible for providing $b_i$ and the upstream utility power line provides $(p_i - b_i)$ . The peak (nameplate) power demand of each rack is $\mathrm{P_r}$ but the allowed maximum power budget $P_{PDU}$ of the PDU is often less than the total peak power $nP_r$ of all the connected racks. To avoid overloading, modern intelligent power distribution unit (iPDU) is able to specify the maximum power of each power outlet.

The power distribution unit (PDU) designates the power budget for each rack and affects the usage of batteries. Given the scaling factor $[\lambda_1, \dots, \lambda_n]$, each power delivery path *i* can assign a maximum power flow (soft limit) of $\lambda_i P_i$. To avoid overloading, the data center must ensure:

$$\mathrm{p_i - b_i \le \lambda_i P_r} \quad (1)$$

$$\mathrm{\sum \lambda_i P_r \le P_{PDU} \le nP_r} \quad (2)$$

Aggressive power provisioning can result in frequent battery usage. Some battery units may incur very low levels of stored energy due to uneven battery discharge. In Figure 5 we present the standard deviation of remaining capacity of 20 rack-mounted batteries at each timestamp. These DEBs are managed in two ways: *offline charging* which recharges whenever the battery capacity drops to a preset threshold; *online charging*, which opportunistically recharges whenever there is additional power budget available. For online charging, the evaluated data center yields roughly 3~12% variation in capacity. Without timely recharge, the *offline charging* nearly doubles the variation in many cases. These aggressively discharged battery units can be extremely vulnerable to power anomalies.

## III. Threat Model and Attack Analysis

DEB systems are the final line of defense against malicious power attacks in most data centers. Their vulnerability demands increasing attention from both data center designers and operators. In this section we specify the types of threats that our system defends against.

### *A. Two-Phase Attack*

Our threat model assumes that a sophisticated adversary can manipulate the power demand of a small group of compute nodes to overload a larger cluster. We assume a two-phase attack tailored to the power provisioning method of today's under-provisioned data centers. Specifically, the entire attack process is organized as three steps:

*1) Preparation: Gain Control of Servers*

To overload the server rack and trigger circuit breaker the attacker first needs to subscribe at least one (preferably multiple) physical machine. These machines will become the host of a *Power Virus*. The attacker can either opportunistically look for such a host by repeatedly creating many virtual machines (VM) and monitoring the IP of the VM instance, or keep rebooting a few VMs until they research the same desired location [24]. Once the attacker has successfully gained control of enough nodes, the next thing is to wait for the best time to attack.

*2) Phase-I Attack: Identify Vulnerable Status*

Servers with inadequate stored energy are much easier to overload. The attacker first identifies vulnerable racks by initiating a "*Non-Offending Power Peak*" which can mildly increase the average utilization of the server rack. In most cases, the data center will treat such power peak as normal load fluctuation (i.e., visible peak). This phase represents the latent period of the power attack.

Figure 6 demonstrates this process using a real battery-backed server cluster (detailed in Section 5). In Phase-I, the attacker keeps running workload in order to accelerate battery discharge. These local batteries become temporarily unavailable since most DEB systems choose to disconnect low-power batteries from load for safety reasons. For example, Facebook uses an independent low-voltage disconnect (LVD) device to isolate the battery unit if the sensed terminal voltage drops below 1.75V per cell [2].

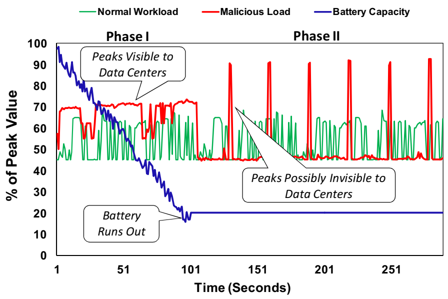

**Fig. 6.** Demonstration of the two-phase attack model

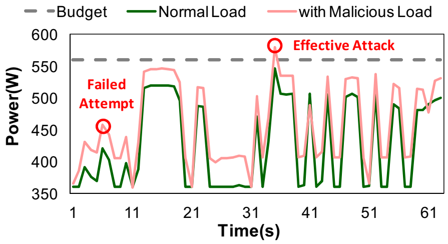

**Fig. 7.** Demonstration of effective power attack

Once the peak-shaving DEB runs out, data center severs have to use performance scaling (e.g., DVFS) to cap power demand. By monitoring the performance of its VMs the attacker would be able to identify when and where the stored energy is low. After multiple times of learning, the attacker can develop the knowledge of the capacity of the associated DEB and estimate the approximate time that the DEB can sustain its "*Non-Offending Power Virus*".

*3) Phase-II Attack: Launch Offending Spikes*

With the above elaborate efforts, the attacker can start to launch "*Offending Hidden Power Spikes*" that will create power spikes possibly invisible to data centers. Before this, the attacker first needs to use the visible peak to drain the battery. Otherwise, these local batteries can eliminate any power quality issues including fine-grained spikes. Afterwards, the attacker can generate short load surges which do not significantly increase the average utilization. As shown in Figure 6, the power virus can be mutated to create very high and narrow power spikes in Phase-II. Most of the existing utilization-based power monitoring mechanism cannot detect such fine-grained power variation [11]. They normally monitor the total energy consumption at coarse-grained intervals (e.g., 10 minutes) to estimate the average power demand. Without enough backup power, the server rack cannot smooth out those power spikes. In this case, the circuit break will be triggered and the service will be temporarily lost, causing catastrophic results.

Data centers today typically lack efficient mechanism to prevent well-planned spikes. Normal power capping mechanisms cannot respond quickly enough to limit the sudden spikes. Advanced power accounting and power capping can operate much faster but it mainly works on per-node level. Oftentimes each server is allowed to reach its peak power as long as the total average rack/PDU utilization is within the budget.

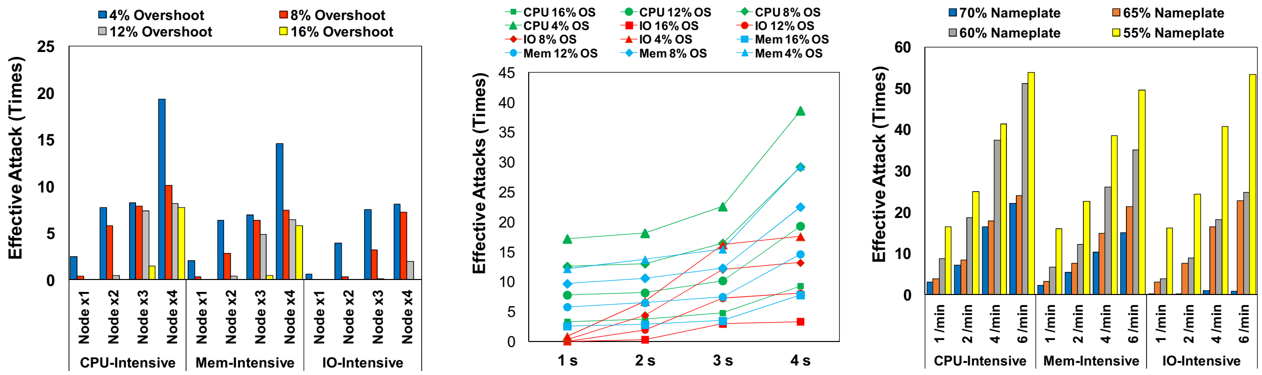

(A) Peak height manipulation (B) Peak width manipulation (C) Attack frequency manipulation

**Fig. 8.** Statistics of effective attacks under various scenarios. In (A), an attacker gains control of different number of server nodes. In (B), the attacker attempts to increase the sustained peak width. In (C), the attacker launches invisible spikes at different frequency. OS: Overshoot; Nameplate: the nameplate power of the system

**Table I. Detection rate under different power metering schemes**

| | | 1 Malicious Server | | | | 4 Malicious Servers | | | |
|---|---|---|---|---|---|---|---|---|---|
| | | W = 1s | | W = 4s | | W = 1s | | W = 4s | |
| | | 1/min | 6/min | 1/min | 6/min | 1/min | 6/min | 1/min | 6/min |
| Metering Interval | 5s | 43.3% | 48.3% | 44.4% | 62.8% | 8.9% | 29.4% | 12.2% | 51.1% |
| | 10s | 37.8% | 45.6% | 43.3% | 68.9% | 1.1% | 14.4% | 16.7% | 100% |
| | 30s | 40.0% | 46.7% | 46.7% | 73.3% | 0.0% | 6.7% | 3.3% | 100% |
| | 60s | 13.3% | 33.3% | 13.3% | 86.7% | 0.0% | 0.0% | 0.0% | 100% |
| | 5m | 0.0% | 0.0% | 0.0% | 100% | 0.0% | 0.0% | 0.0% | 100% |
| | 10m | 0.0% | 0.0% | 0.0% | 100% | 0.0% | 0.0% | 0.0% | 100% |
| | 15m | 0.0% | 0.0% | 0.0% | 100% | 0.0% | 0.0% | 0.0% | 100% |

By creating high and short power spikes on multiple servers, it is very likely to exceed the limit of the circuit breaker. Whether or not an effective attack can trip the circuit breaker depends on the actual over-current and the peak current duration [11]. Tripping a circuit breaker is not an instantaneous event since most PDU can tolerate certain degrees of brief current overloads. However, once the overload exceeds certain threshold, it requires very short time (several seconds) to trip a circuit breaker.

Note that we do not argue that the aforementioned "*Offending Power Virus*" will guarantee a successful attack (i.e., power failure). A single power spike may not necessarily result in effective attack (i.e., power draw exceeds a pre-determined limit), since other normal servers might incur power valley at the same time. Repeatedly creating hidden power spikes could eventually lead to an overload, as shown in Figure 7. Given enough overload events, it has very good chances to fail a server rack.

### *B. Power Attack Analysis*

The way the attacker launches power spikes greatly affects attacking results. Figure 8 presents the impact of peak power manipulation on the number of effective attacks for 15 minutes. We consider three key factors: the height, width, and frequency of power spikes.

Figure 8-A shows the increase in effective attacks under different numbers of malicious nodes. The x% overshoot indicates the maximum power overload that the data center can tolerate. It is clear that gaining control of more machines eases power attack. Particularly for an IO-intensive power virus, the attacker might need more servers to increase the chance of a successful attack.

Meanwhile, the attacker can also accelerate the attacking process by increasing the width of power peaks. We note from Figure 8-B that, increasing peak duration can greatly increase effective attacks at certain point. For example, a 4-second CPU-intensive power virus yields almost 2X effective attacks than a 3-second power virus.

Further, an attacker can launch even more aggressive attacks by generating frequent power spikes. In Figure 8-C we consider different rates from 1 to 6 times per minute for a 1-second CPU-intensive power virus. It shows that there is a positive correlation between attack frequency and effective attack number. But the latter is not in proportional with the former. In addition, since the I/O intensive power virus cannot effectively trigger high spikes in the Phase II, it may fail to create any effectiveness attack when the power budget is adequate (e.g., 70% nameplate power).

As the attacker uses more aggressive attack approaches (increasing spike duration, frequency, etc.), the chances of being able to be detected by the data center also increase. Advanced power metering and complex power management software allow for higher detection rate, but not all data center can afford the significant overhead of such fine-grained profiling [11]. We evaluate the detection rate of various power attacking scenarios under different power demand monitoring technologies for 15 minutes. In Table 1 we show that even fine-grained power monitoring cannot detect all the hidden power spikes. For example, only 40~60% power spikes are detected using a power meter that measures load power demand every 5 seconds. In many cases, the data center is totally blind to fine-grained power spikes. Although the data center can apply cluster-wide power capping to eliminate any hidden power spikes, such security measures may well be overkill and could significantly affect other legitimate service requests.

# IV. Power Attack Defense

We propose power attack defense (PAD), a new design patch that allows emerging battery-backed data centers to run safely and smoothly under power-related attacks. It has three distinctive features:

- *Hierarchical Emergency Handling.* PAD defines the security policy for battery-backed server clusters. It lays down the general rule for protecting data centers from different types of malicious loads that intend to overload the system.
- *Joint Software-Hardware Support.* PAD creates a virtual energy backup pool, called virtual DEB (vDEB), to increase attacking cost. It also introduces a dedicated energy backup device, called micro DEB (μDEB), to handle undetectable spikes.
- *Heterogeneous and Hybrid Backup.* PAD uses a heterogeneous DEB architecture that combines fine-grained and coarse-grained integration. It leverages both batteries and emerging super-capacitors to provide better energy/power support.

## A. Basic Management Policies

PAD adopts a hierarchical model, where power management strategies are classified into different levels of emergency states. We have defined three levels: Normal (Level 1), Minor Incident (Level 2), and Emergency (Level 3). There are three inputs that affect the state: vDEB, μDEB, and VP that indicates if a visible peak is identified. Figure 9 illustrates different states and the flows between them.

Our policy defines the initial states for all the combinations of initial inputs. Depending on the operating environment, PAD may enforce different security levels and expose underlying power/energy profile to data center. This allows the data center to make informed decision on secure power management. For example, if the data center undergoes sustained power peaks (i.e., visible peaks) in Level 1, it will intelligently enable a fraction of DEB units to shave the power peak (detailed in Section 5.2). In contrast, if PAD believes that the data center is under the threat of potential hidden spikes in Level 2, it will keep a watchful eye on the health of the μDEB and collect load information for future inspection and anomaly prevention.

In rare cases, when both vDEB and μDEB are empty, PAD will overlook the load power behavior and force to enter an emergency state. This can cause the data center to shed loads, i.e., put some servers into sleeping/hibernating states or trigger load migration from vulnerable racks to dependable racks. Although the temporary load shedding may incur certain performance degradation, it is not overkill. This could prevent data center from incurring significant loss during a large-scale power failure. In fact, by sleeping only a small amount of servers, one can prevent the majority of data center racks from power-related attacks.

Note that the initial level for [vDEB>0, μDEB <0] is not specified. This is because it is not a stable energy backup state since the μDEB can always be charged by vDEB which has much larger energy capacity. As shown in Figure 9, one can use either Level 1 or Level 2, depending on the level of security requirement of the organization.

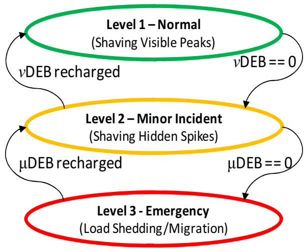

| vDEB >0? | μDEB >0? | VP >0? | Initial State |
|---|---|---|---|
| 0 | 0 | 0 | L3 |
| 0 | 0 | 1 | L3 |
| 0 | 1 | 0 | L2 |
| 0 | 1 | 1 | L3 |
| 1 | 0 | 0 | (L1/L2) |
| 1 | 0 | 1 | (L1/L2) |
| 1 | 1 | 0 | L1 |
| 1 | 1 | 1 | L1 |

**Fig. 9.** Hierarchical security level defined by PAD. The initial state is determined based on the monitored peak power information (VP>0 means a visible peak power is detected) and the available backup energy in the virtual DEB and micro DEB

## B. PAD Architecture

The driving insight of our work is that the main source of vulnerability lies in *the reliance on a traditionally very simple, homogeneous DEB architecture to defend against a potentially variable and sophisticated power anomaly*. To tackle the security challenge faced by exiting data centers, PAD has adopted several important mechanisms. Figure 10 shows the schematic diagram of the PAD design.

In the following we first discuss the PAD's virtual DEB design, which aims at protecting a Level-1 data center from a brute visible peak attack. We then discuss micro DEB design, which intend to defend against a more sophisticated hidden spike attack often seen in a Level-2 state.

### 1) Virtual Distributed Energy Backup (vDEB)

Rather than treating rack-mounted batteries as separated energy backup systems, PAD creates a virtual energy backup pool termed vDEB and a vDEB controller for managing it. The vDEB controller, which is enlightened by the power capacity sharing mechanism at the PDU level, allows vulnerable server racks to share unused energy backup capacity within the same PDU.

Conventionally, server rack power allocation is largely workload-driven and consequently overlooks the pressure the server rack may exert on batteries. In addition, recent battery-based peak power management schemes are largely battery lifetime-driven and fail to consider the uneven usage of batteries. Consequently, some racks may aggressively discharge their batteries and at some point happen to become the weak point of data center. Once a PDU level power failure occurs, each server rack will become a standalone system that can only draw power from its local energy backup. If the autonomy time (the maximal outage duration that the battery can support) is not long enough, the data stored on the server can be damaged.

Our vDEB controller uses an intelligent algorithm for managing uneven battery usage. It combines cluster-level battery balancing and rack-level battery balancing. We assign the discharge rate of each battery unit based on the available SOC value (Algorithm 1). This prevents vulnerable

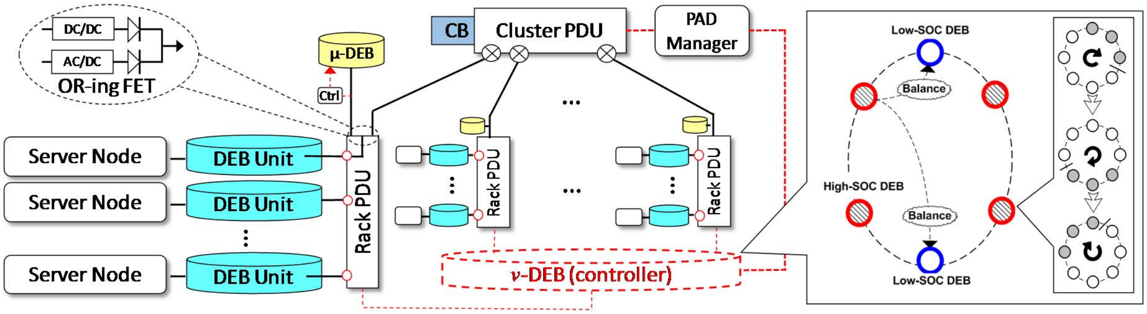

**Fig. 10.** The PAD architecture and power management scheme

**Algorithm 1**: Heuristic for the *v*DEB two-level load sharing

```
socList [$s_1, s_2, \ldots, s_n$] ← GetCurrentSOC();
powerList [$p_1, p_2, \ldots, p_n$] ← GetCurrentPower();
$P_{total} \leftarrow \sum powerList$
$SOC_{total} \leftarrow \sum socList$
$P_{shave} \leftarrow P_{total} - P_{max}$;
if ($P_{shave} > P_{ideal}$)
    evenly usage DEB
else
    Quicksort rack ID based on the SOC value of each rack;
    Store the quick sort (descending) in R[i];
    for (i=1; socList[R[i]]/ $SOC_{total}$ * $P_{shave}$) > $P_{ideal}$ && i< N; i++)
        P[R[i]] ← $P_{ideal}$;
        $SOC_{total} \leftarrow SOC_{total}$ – socList[R[i]];
        $P_{shave} \leftarrow P_{shave} - P_{ideal}$ / N;
    end
    for (j=i; j<N; j++)
        P[R[j]] ← SOC[R[j]]/$SOC_{total}$ * $P_{shave}$;
    end
end
Discharge DEB based on P[i] in each rack in round-robin
```

batteries from aggressively discharging and allows for fast balancing. In the meantime, we also notice that batteries have a maximum discharge rate for reliability and safety reasons (e.g., normally 48A for a 2Ah lead-acid battery cell) [25]. Therefore the discharge algorithm should not cause accelerated aging on battery systems. We have set an upper bound when assigning the discharge rate (i.e. represented by the ideal discharge power $P_{ideal}$).

The vDEB design brings two important benefits. First, it allows a data center to hide a vulnerable battery-backed server rack. It greatly extends the peak shaving time during a Level-1 power management process. As a result, the cost of bringing down a server can increase significantly. On the other hand, vDEB can often frustrate an attacker's efforts to gain critical information such as "*how long does the victim rack's battery can sustain*". This is because the capacity sharing mechanism involves multiple server racks that an attacker may not gain access to (adding considerable noise to an attacker's observations in a side-channel attack).

*2) Micro Distributed Energy Backup (μDEB)*

Virtual DEB alone cannot defeat a well-planned power attack. A power virus can be mutated to create transient power spikes that most utilization-based power management software cannot detect. As a result, one cannot timely enable the server-level DEBs to provide the necessary backup energy support.

We propose to further integrate a dedicated small power backup device in existing rack power zone to existing distributed battery system (Figure 10). The device, termed as micro DEB (μDEB), is designed to further strengthen the defense against hidden power spikes at the server rack level. In order to protect server racks from undetectable power anomalies, the μDEB must be designed to react to any voltage surge/sags automatically. To this end, we connect μDEB with the primary power delivery bus using an ORing controller (a low forward-voltage FET device), as shown in Figure 10. The ORing has been widely used in today's redundant power sources to enable hot swaps and current sharing. In this study we leverage it to design a spike-shaving system. This idea does not apply to peak-shaving for two reasons. First, at the server level, current sharing can result in degraded efficiency in server power supply unit. Second, at the rack level, current sharing for sustained peak shaving can cause thermal issues in μDEB.

Shaving the transient power spike requires very small energy capacity but very large power output capability. This motivates us to use the promising super-capacitor (SC) system instead of conventional lead-acid battery. SC is expensive (10~30$/Wh) but μDEB does not require very large capacity. For example a 5KW power rack for 0.5 second current sharing only requires 0.35Wh backup energy capacity. Normally a 2A battery cell can provide 85 W at the maximum for 5minutes [25]. This requires us to connect many battery cells in parallel to achieve the desired power capacity, which can be bulky and expensive. Further increasing the output power requires higher output current which can greatly accelerate the aging of lead-acid batteries [27]. In contrast, super-capacitors can provide very impressive power output with no lifetime concerns. It has much smaller footprint and is environmental-friendly too.

The dedicated hardware component of PAD is necessary since it defends against invisible power spikes. Although some software mechanisms such as Intel's RAPL allows a data center to better monitor and manage load power, one cannot completely rely on them to handle malicious loads. This is because predicting total server power demand (not just CPU/Mem) quickly and accurately is still challenging.

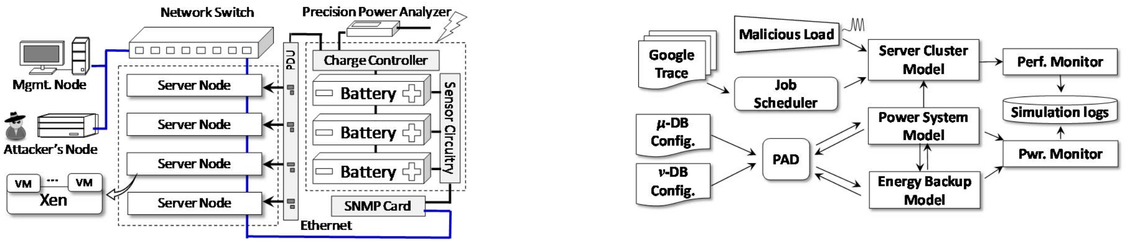

(A) Scaled-down hardware system (B) Trace-based simulation framework

**Fig. 11.** The validation and evaluation platform of PAD

**Table II. Evaluated Attack Scenarios**

| Schemes | Descriptions |
|---|---|
| CPU Intensive | Threaded Tachyon, a parallel ray-tracking system [28] |
| Mem Intensive | Stream, a system memory testing program [29] |
| IO Intensive | Apache benchmark with 1 million requests [30] |

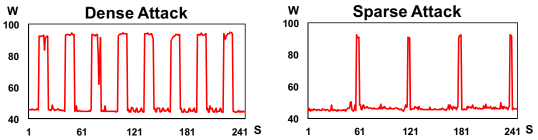

**Fig. 12.** Example of the collected attacking traces. Left: dense and extensive attack. Right: sparse and light-weighted attack

**Table III. Evaluated Power Management Schemes**

| Schemes | Descriptions |
|---|---|
| *Conv* | Conventional designs that do not discharge batteries dynamically and only use them to handle outage |
| *PS* | Recent peak shaving schemes that use energy backup in each BBU to handle visible power spikes |
| *PSPC* | Combing PS with power capping mechanism which can decrease processor frequency by 20% |
| *vDEB* | vDEB-only design: PS + load sharing mechanism that can eliminate vulnerable racks. |
| *μDEB* | μDEB-only design: PS + micro energy backup devices that can handle the rack-level power spikes. |
| *PAD* | Our power management patch for securing data center from both visible and hidden power attack. |

If datacenters perform power capping based on inaccurate power monitoring, they can cause significantly degraded performance on normal loads. Even if full-system accurate power prediction is available, it often takes 100ms~300ms to reduce the power demand, which is not fast enough to correctly shave the peak under the rapid power dynamics observed in data centers [26]. As long as the load current and attacking time are well controlled by the attacker, the PDU circuit breaker can still be tripped [11]. Furthermore, the cost of per-server proactive monitoring with fine granularity can be more expensive than hardware-based peak power shaving. In fact, even the Top500/Green500 HPC data centers only sample power at one-second intervals for ranking purposes.

## V. Evaluation Methodology

We build a scaled-down testing platform as shown in Figure 11-A. It consists of a mini server rack and a set of three YUASA UPS batteries. The total power capacity is 800W and it can maintain 10 minutes under full load. All the batteries are dynamically monitored on a per minute basis. Our system is able to dynamically switch ON/OFF the UPS with SNMP commands over Ethernet and collect key battery and power information during runtime.

We model different power viruses that take advantage of three types of benchmarks: CPU, Memory, and I/O intensive, as shown in Table 2. We deploy the benchmark on Ubuntu (14.01 LTS) virtual machines created on Xen 6.5.0 hypervisor. We create power virus on our hardware platform and collect the power activity trace of our system using a precision power meter that has a maximum sampling rate of 200KS/s and less than 0.1% error rate. Figure 12 shows power virus trace examples we generated. Based on the configuration of our system, we consider two types of power attack: a dense and extensive power spikes and a sparse and less aggressive spikes

We feed the collected power virus traces to a trace-based data center simulator that takes real Google compute traces [15] as input (Figure 11-B). The Google cluster trace (with an interval of 5 minutes) represents 1-month worth of node information from May 2010, on a cluster of about 220 machines. Work arrives at the cluster in the form of jobs. A job is comprised of one or more tasks, each of which is accompanied by a set of resource requirements used for dispatching the tasks onto machines. Every line in this trace includes start time, end time, machine ID, and CPU rate of the task. We use programs to process the trace in our event-driven simulation platform. We use machine ID as the identifier and calculate the total CPU power demand belong to a given machine at the same timestamp.

We assume a HP high-performance ProLiant DL585 G5 server system (2.70GHz, AMD Opteron 8384), which has an active idle power of 299W and a peak power of 521W [31]. There are 22 racks in total and each rack has 10 servers. In this work we assume a DEB system similar to Facebook's V1 design [2]. Each rack has a dedicated battery cabinet for power shaving. The fully charge battery can sustain 50 seconds under full load. We maintain detailed charge/discharge logs and calculate the capacity decrease and increase using a kinetic battery model (KiBaM) [32] at each fine-grained timestamp throughout the simulation. All the power system models are embedded in our simulation platform as shown in Figure 11-B.

Table 3 summarizes our evaluated six power management schemes for battery-backed data center. We consider three

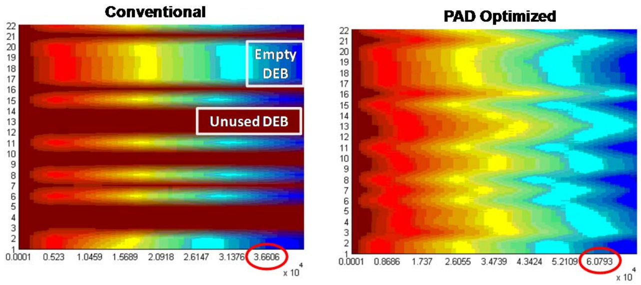

**Fig. 13.** A comparison of DEB usage in conventional datacenters and datacenters protected by PAD. (x-axis: seconds; y-axis: rack ID)

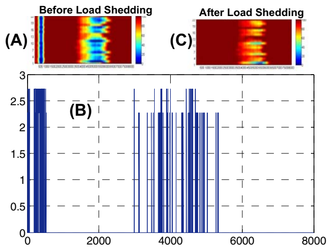

**Fig. 14.** Shedding less than 3% load could avoid aggressive battery usage

baseline data centers: *Conv*, *PS*, and *PSPC*. *Conv* represents the most traditional data center designs that only use centralized battery as power backup. *PS* uses the state-of-the-art power shaving schemes similar to [7]. *PSPC* further combines *PS* with performance scaling mechanisms (i.e., DVFS) for better design trade-offs. It aims at efficiently shaving the peak power but overlooks the power-related security issue of DEB-based data centers. We have compared the above baseline schemes with a vDEB-only design, a µDEB-only design, and finally PAD that combines vDEB and µDEB.

## VI. Evaluation Results

In this section we first present detailed DEB profiling map to illustrate the optimization effectiveness of PAD. We then evaluate the survival time (i.e., from the beginning of the attack to the time the first overload happens) of different power management schemes under various attack scenarios. Afterwards we show that PAD provides better performance guarantee and incurs minor cost overhead.

### A. *Effectiveness*

We examine the energy backup usage profile in data centers over a one day period. It is clear that vulnerable racks exist in conventional data centers. Figure 13 shows the monitored DEB utilization map of the evaluated server clusters at each timestamp. In the figure, dark red represents fully charged batteries while dark blue means near-empty batteries and the associated vulnerable racks which could be ideal targets for a sophisticated criminal. We note that some server racks in conventional data centers have to heavily discharge their associated DEB systems to reduce peak power demand. The battery utilization pattern in this case becomes highly dependent on the power behavior of each individual rack and therefore exhibits significant variation.

Our result shows that PAD allows a data center to hide vulnerable server racks by effectively balancing the usage of batteries. Although uneven usage still exists to some extent, those server racks no longer differ significantly in their backup power at any timestamp. As a result, the survival time is improved by 1.7X after optimization. We recognize that PAD cannot completely eliminate overload under constant aggressive attack. Our main objective is to extend the sustained operation time as much as possible to frustrate the attacker's plan by significantly increasing the cost of launching a successful attack. In addition, it also gives operators more time to identify malicious loads and figure out any possible solutions. Figure-13 evaluates a small cluster and therefore the results are not striking. In data centers that have hundreds/thousands of racks, PAD can offer impressive security/availability benefits.

Note that PAD never uses aggressive server shedding to save battery energy. PAD temporarily put some of the low-priority racks into deep-sleep mode only in extreme cases when cluster-wide power peaks appear. This has two major benefits. First, it prevents potential data corruption in the event of an unexpected overload. Second, shutting down some vulnerable loads may disrupt the attack process. The insight here is that we can actually eliminate the power shortfall and release the burden on the DEB system by just shutting down a very small amount of non-critical loads. In Figure 14 we investigate a periodic data center-wide load surge that can create massive amount of vulnerable racks in conventional designs (Figure 14-A). Our result shows that a load shedding ratio of about 3% of the entire data center servers (Figure 14-B) is able to achieve an impressive balanced battery usage map (Figure 14-C). In Figure 14-C, PAD has successfully avoided the narrow blue strip and mitigated rack vulnerability of the wide blue strip.

### B. *Survival Time*

We further quantify the security benefits of PAD. We focus our attention on how long the data center will sustain under power attack. Figure 15 shows the evaluation results across different power virus scenarios.

Our first observation is that the way how an attacker generates power virus can affect the attacking results. In Figure 15, the CPU-intensive power virus is more likely to trigger effective power spikes, and therefore result in lower sustained time. Although the results of conventional design *Conv* is not sensitive to power virus, PAD shows remarkable survival time improvement under a light-weight attacking (i.e., sparse attack + I/O workload)

Figure 15 also demonstrates the different impacts of *µDEB* and *vDEB* on data center survival time. We find that both *µDEB* and *vDEB* could extend server survival time but the improvements of *vDEB* are bigger. This mainly because that the visible power peaks dominate in the overall attacking period. Combing *µDEB* and *vDEB* allows PAD to better deal with power virus. Overall, PAD improves the sustained time

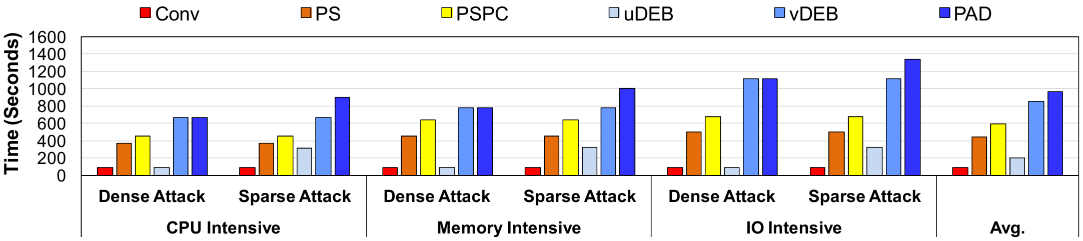

**Fig. 15.** The sustained operation duration of the evaluated Google cluster under various power attacks

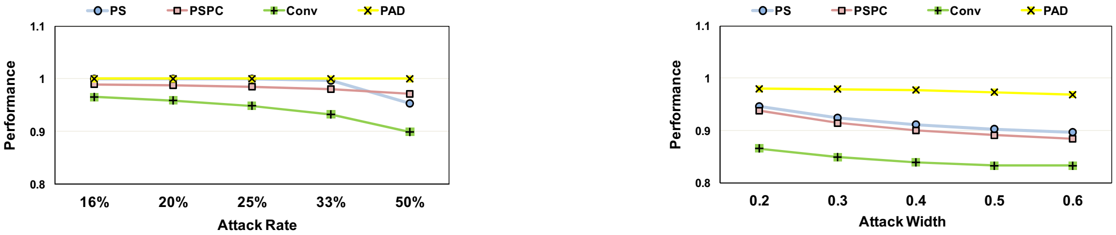

(A) Different power attack rate

(B) Different power attack width

**Fig. 16.** The overall data center throughput during the attack period

by 10.7X compared to conventional data centers, and 1.6X compared to the state-of-the-art proposals.

## C. *Performance*

It is often equally important that a security-aware design does not compromise performance. In fact, since PAD is orthogonal to existing system and software level optimizations, it has no side-effects on workload performance during normal operation. On the contrary, because PAD can greatly intend the sustained operation time under power attack, it can greatly reduce unnecessary power capping activities that are seen in other baselines.

We evaluate the total data center throughput under different power attack rates and peak power widths. We note from Figures 16-A and 16-B that, as power attack becomes more aggressive, existing schemes can result in performance degradation to some extent. For example, the throughput can drop by roughly 10% when we increase the attack rate to 50%. Compared to attack rate, the peak power width has more notable performance impact. Even so, PAD shows less than 5% throughput degradation for the evaluated 0.6s power spike, while the performance degradation of *PSPC* and *Conv* are 12% and 17%, respectively.

## D. *Cost Efficiency*

The major hardware addition in our design is μDEB which uses small-scale super-capacitors to shield data center from invisible power spikes for many times. In this work vDEB is not treated as cost overhead since we leverage battery devices that most data centers already have.

PAD incurs minor cost overhead. In Figure 17 we examine the relationship between cost overhead and the survival time. The cost of μDEB mainly depends on its capacity, which roughly follows a linear model. One can keep the cost of μDEB below certain percentage of vDEB by limiting the installed capacity of μDEB.

Importantly, our result implies that a small increase in μDEB capacity can have a large impact on the sustained time of PAD. As shown in Figure 17, increasing the capacity of μDEB from 1% to 15% could extends the data center emergency handling capability (i.e., survival time) by nearly 40X. Although it is evident that a larger μDEB increases the emergency handling capability, the associated cost also mounts up. We expect that companies will adopt different capacity planning strategies to achieve their desired trade-offs in profitability, availability, and security.

# VII. Related Work

In this section we discuss representative prior studies in different domains that are most relevant to our work.

## A. *Power-Related Attack*

Very little prior work exists on enhancing the data center's vulnerability to power-related attacks. Recently, Xu et al. [6] demonstrate the feasibility of launching simultaneous power spikes on servers to trigger outage but does not consider the energy backup. We argue that such an attack is more likely to happen when the energy backup system is vulnerable. In addition, the authors only focus on conventional power capping and server consolidation to mitigate the negative impact of malicious loads. In contrast, we look at DEB system which represents a new paradigm of data center power management and control.

Many prior works are focused on the attack against electric power grid [33-35], which could also result in data center power outage indirectly. For example, Liu et al. [33] show how to compromise today's complex utility power grid system. Soltan et al. [34] explore the cascading failure of the utility grid power transmission line. Chen et al. [35] propose to leverage energy storage to mitigate cascading failures of utility grid. In general, all these works focus on attack through false data/code injection. Differently, we investigate

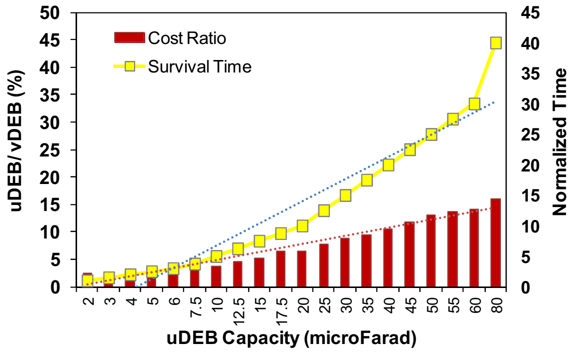

**Fig. 17.** Cost efficiency analysis

a newly merged attack approach which intends to overload servers that have limited backup support.

### B. *Secure Architecture*

Building secure architecture is drawing growing attention. Many prior works are focusing on various side-channel attacks. For example, Demme et al. [36] propose a metric for measuring information leakage called side-channel vulnerability factor (SVF). Colp et al. investigate the data protection issue on mobile devices [37]. Chen et al. propose a microarchitecture-level framework to detect the possible presence of covet timing channels [38]. Liu et al. introduce a new classification of cache side channel attacks [39]. Callan et al. propose a practical methodology for measuring the side-channel signal [40]. In contrast to these work, we focus on a new type of DDOS-like threat that intend to crash or interrupt server rack operation by abusing its power and energy resources.

### C. *Data Center Availability*

Another representative group of related work is in the context of data centers availability and reliability.

There have been prior works on the reliability of data center infrastructure. For example, Govindan et al. [41] model the reliability of data center power infrastructure using Markov Chain and Reliability Block Diagrams. This paper shows that hybrid UPS system could improve the overall reliability but lacks a detailed scheduling scheme to manage the system. Sankar et al. [42] mainly investigate power redundancy related design trade-offs, while our work investigates the vulnerability in backup power systems. On the other hand, cooling system could also become a vulnerable part in data centers [43, 44]. Compared to data center brownout due to thermal runaway, server outage caused by overloading can be much easier to happen.

Many research works are centered on the reliability analysis of computing/networking systems. For example, Vishwanath et al. [45] have characterized the hardware reliability in cloud data centers. Sankar et al. have investigated the soft failures in data centers [46] and have established the correlation between temperatures and hard disk drive failures [47]. Wu et al. [48] propose to mitigate network failures by deactivating suspected components. Mysore et al. [49] propose a fault-tolerant network infrastructure that is aware of the baseline topology. In [50], researchers from Microsoft thoroughly investigate the network failures in a data center. All these works are mainly focused on the component/device failure. Our work differs from them in that we focus on protecting servers from power attack launched by a sophisticated adversary.

### D. *Battery Management*

Conventionally, batteries are only used as emergency backup which is rarely used. Recently, batteries have been used to shave the occasional load power peaks [8, 9, 13]. For example, Govindan et al [8] propose the concept of energy buffer (eBuff) in data centers. Wang et al. [9] provide a detailed characterization of various energy storage technologies. However, these prior works are only limited to the dynamic scheduling and cost optimization of energy storage systems. In [13], Govindan et al. propose to use UPS batteries to avoid costly power capping or defer virtual machine migration. This work mainly focus on using static battery management to handle normal load power emergency, while our work explores dynamically coordinating massive battery units to handle attacks.

On the other hand, batteries are also used in emerging green data centers to temporarily store the excess renewable energy generation or handle the power shortfall when renewable energy is inadequate [15-17, 51, 52]. Normally, batteries can be placed both at the data center level [15, 16] or the server level [17]. They are mainly used to improve green energy utilization and avoid service interruption. Differently, in this work we focused our attention on the security and availability of battery-dependent data centers.

A few recent proposals have focused on deploying and managing distributed batteries [7, 53, 54]. For example, Kontorinis et al. [7] explore distributed UPS systems for power capping in data centers. Aksanli et al. [53] optimize the efficiency of peak power shaving. Ghai et al. [54] have proposed a controller for distributed local energy storage devices to optimize power supply efficiency. In [4], researchers have compared the efficiency of different power supply systems that use distributed batteries.

Although some of the existing techniques do exploit ad-hoc discharge control for extending battery lifetime [7, 8, 15, 17], they do not consider the risks that aggressive battery usage may have and fail to timely eliminate vulnerable batteries. Recent works have explored the usage of hierarchical and hybrid DEB in data centers, but they only focus on energy efficiency and does not consider the security issue of aggressive power management [55, 56]. As a result, the associated servers are often the potential victims of power virus. Prior researches largely overlook this issue and this paper aims to fill this critical void.

**Summary of Novelty:** The novelty of this paper is three -folded. (1) It is the first to explore the vulnerability of emerging DEB-backed data centers under power virus that can generate visible peaks and hidden spikes. (2) It defines a new security policy for DEB-based data center and lays down the general rule for protecting data centers from malicious loads. (3) It proposes a novel data center design patch called PAD to shield servers from power-related attack using light-weight software/hardware co-designs.

## VIII. Conclusions

Driven by energy-efficiency and cost, future large-scale computing infrastructure is projected to be backed by massive small-scale distributed energy backup (DEB) rather than a central UPS system. To safely exploit the benefits of distributed batteries that distributed energy storage units may provide, data center designers need to understand the security issue of these systems. In this paper we propose a security-based power management for mitigating the system's vulnerability to power attacks. The proposed design allows data centers to smartly plan their usage of DEB units and enables the servers to operate smoothly for extended duration (1.6~11X) with better performance guarantee and negligible cost overhead.

## ACKNOWLEDGEMENT

We thank all the reviewers for their valuable comments. This work is supported, in part, by the National Natural Science Foundation of China (No. 61502302, 61202026 and 61332001), Program of China National 1000 Young Talent Plan, and the National Basic Research Program of China (973 Program, No. 2015CB352403).

## References

[1] Google uncloaks once-secret server, 2009 http://www.cnet.com/news/google-uncloaks-once-secret-server-10209580/

[2] P. Sarti. Battery Cabinet Hardware v1.0, Open Compute Project, 2012. http://www.opencompute.org/

[3] Microsoft Reinvents Datacenter Power Backup with New Open Compute Project Specification, 2015. http://blogs.technet.com/b/msdatacenters/archive/2015/03/10/microsoft-reinvents-datacenter-power-backup-with-new-open-compute-project-specification.aspx

[4] Y. Kuroda, A. Akai, T. Kato, and Y. Kudo. High-Efficiency Power Supply System for Server Machines in Data Center, *International Conference on High Performance Computing and Simulation* (HPCS), 2013

[5] HP Flexible Slot Power Supplies, http://www8.hp.com/us/en/products/power-supplies/product-detail.html?oid=7268787

[6] QuantaPlex T21SR-2U Datasheet, http://www.quantaqct.com/

[7] V. Kontorinis, L. Zhang, B. Aksanli, J. Sampson, H. Homayoun, E. Pettis, T. Rosing and D. Tullsen, Managing Distributed UPS Energy for Effective Power Capping in Data Centers, *International Symposium on Computer Architecture* (ISCA), 2012

[8] S. Govindan, A. Sivasubramaniam and B. Urgaonkar. Benefits and Limitations of Tapping into Stored Energy for Datacenters, *Int. Symp. on Computer Architecture* (ISCA), 2011

[9] D. Wang, C. Ren, A. Sivasubramaniam, B. Urgaonkar, and H. Fathy. Energy Storage in Datacenters: What, Where, and How Much, *SIGMETRICS Performance Evaluation Review, Vol. 40, No. 1*, 2012

[10] Z. Xu, H. Wang, Z. Xu, and X. Wang. Power Attack: An Increasing Threat to Data Centers, *The Network and Distributed System Security Symposium* (NDSS), 2015

[11] D. Meisner, and T. Wenisch Peak Power Modeling for Data Center Servers with Switched-Mode Power Supplies, *International Conference on Low Power Electronic Design* (ISLPED),2010

[12] X. Fan, W. Weber, and L. Barroso, Power Provisioning for a Warehouse-Sized Computer*, International Symposium on Computer Architecture* (ISCA), 2007

[13] S. Govindan, D. Wang, A. Sivasubramaniam, and B. Urgaonkar. Leveraging Stored Energy for Handling Power Emergencies in Aggressively Provisioned Datacenters. *International Conference on Architectural Support for Programming Languages and Operating Systems* (ASPLOS), 2012

[14] Í. Goiri, R. Beauchea, K. Le, T. Nguyen, M. Haque, J. Guitart, J. Torres, and R. Bianchini. GreenSlot: Scheduling Energy Consumption in Green Datacenters, *Supercomputing* (SC), 2011

[15] I. Goiri, W. Katsak, K. Le, T. Nguyen, and R. Bianchini. Parasol and GreenSwitch: Managing Datacenters Powered by Renewable Energy, *International Conference on Architectural Support for Programming Languages and Operating Systems* (ASPLOS)*,* 2013

[16] C. Li, A. Qouneh, and T. Li. iSwitch: Coordinating and Optimizing Renewable Energy Powered Server Clusters, *International Symposium on Computer Architecture* (ISCA), 2012

[17] C. Li, Y. Hu, R. Zhou, M. Liu, L. Liu, J. Yuan, and T. Li. Oasis: Enabling Datacenter to Scale Out Economically and Sustainably, *International Symposium on Microarchitecture* (MICRO), 2013

[18] Ponemon Institute. 2013 Study on Data Center Outages

[19] Ponemon Institute. 2013 Cost of Data Center Outages

[20] J. Williams, Data Center Security Survey. SANS Institute, 2014

[21] Google Trace. https://code.google.com/p/googleclusterdata/

[22] W. Turner and K. Brill, Cost Model: Dollars per kW plus Dollars per Square Foot of Computer Floor. *White Paper*. Uptime Institute, 2008

[23] Understanding Electric Demand. *White Paper*, National Grid, 2012. https://www.nationalgridus.com/niagaramohawk/non_html/eff_elec-demand.pdf

[24] T. Ristenpart, E. Tromer, H. Shacham, and S. Savage. Hey, You, Get Off of My Cloud: Exploring Information Leakage in Third-Party Compute Clouds. *The ACM Conference on Computer and Communications Security* (CCS), 2009

[25] 12v 12Ah Lead Acid Battery, http://www.micropik.com/PDF/CP12120.pdf

[26] A. Bhattacharya, D. Culler, A. Kansal, S. Govindan, S. Sankar. The need for speed and stability in data center power

capping, *Sustainable Computing: Informatics an Systems, Vol. 3, Issue 3, pp. 183-193*, 2012

[27] L. Liu, C. Li, H. Sun, Y. Hu, J. Gu, and T. Li. BAAT: Towards Dynamically Managing Battery Aging in Green Datacenters, *Annual IEEE/IFIP International Conference on Dependable Systems and Networks* (DSN), 2015

[28] http://jedi.ks.uiuc.edu/~johns/raytracer

[29] https://openbenchmarking.org/test/pts/stream

[30] https://openbenchmarking.org/test/pts/apache

[31] SPECpower_ssj2008 Results.

http://www.spec.org/power_ssj2008/results/

[32] M. Jongerden and B. Haverkort, Which Battery Model to Use?, In *Special Issue on Performance Engineering*, 2009

[33] Y. Liu, P. Ning, and M. Reiter. False Data Injection Attacks against State Estimation in Electric Power Grids, *ACM Conference on Computer and Communications Security* (CCS), 2009.

[34] S. Soltan, D. Mazauric, and G. Zussman. Cascading Failures in Power Grids - Analysis and Algorithms, *ACM e-Energy*, 2014

[35] X. Chen, W. Yu, D. Griffith, N. Golmie, and G. Xu. On Cascading Failures and Countermeasures based on Energy Storage in the Smart Grid, *IEEE Real-Time and Embedded Technology and Applications Symposium* (RTAS), 2014

[36] J. Demme, R. Martin, A. Waksman, and S. Sethumadhavan. Side-channel Vulnerability Factor: A Metric for Measuring Information Leakage. *International Symposium on Computer Architecture* (ISCA), 2012

[37] J. Chen and G. Venkataramani. CC-Hunter: Uncovering Covert Timing Channels on Shared Processor Hardware, *International Symposium on Microarchitecture* (MICRO), 2014

[38] F. Liu and R. Lee. Random Fill Cache Architecture. *International Symposium on Computer Architecture* (ISCA), 2014

[39] P. Colp, J. Zhang, J. Gleeson, S. Suneja, E. Lara, H. Raj, S. Saroiu, and A. Wolman. Protecting Data on Smartphones and Tablets from Memory Attacks. *International Conf. on Architectural Support for Programming Languages and Operating Systems* (ASPLOS)*,* 2015

[40] R. Callan, A. Zajic, and M. Prvulovic. A Practical Methodology for Measuring the Side-Channel Signal Available to the Attacker for Instruction-Level Events, *International Symposium on Microarchitecture* (MICRO), 2014

[41] S. Govindan, D. Wang, L. Chen, A. Sivasubramaniam, and B. Urgaonkar. Towards Realizing a Low Cost and Highly Available Datacenter Power Infrastructure, HotPower, 2011

[42] S. Sankar, D. Gauthier, S. Gurumurthi. Power Award Provisioning in Large Data Centers, *ACM International Conference on Computing Frontiers* (CF), 2014

[43] R. Zhou, Z. Wang, C. Bash, T. Cade, and A. McReynolds. Failure Resistant Data Center Cooling Control Through Model-Based Thermal Zone Mapping, *Technical Report. HP*, 2012. http://www.hpl.hp.com/techreports/2012/HPL-2012-69.pdf

[44] S. Shields. Dynamic Thermal Response of the Data Center to Cooling Loss During Facility Power Failure. *Master Thesis, Georgia Institute of Technology*, 2012

[45] K. Vishwanath and N. Nagappan. Characterizing Cloud Computing Hardware Reliability. *ACM Symposium on Cloud Computing* (SoCC), 2010

[46] S. Sankar and S. Gurumurthi. Soft Failures in Large Datacenters, *IEEE Computer Architecture Letters, Vol 13, NO. 2,* 2014

[47] S. Sankar, M. Shaw, K. Vaid, and S. Gurumurthi. Datacenter Scale Evaluation of the Impact of Temperature on Hard Disk Drive Failures. *ACM Trans. on Storage, Vol. 9, No. 2, Article 6*, 2013

[48] X. Wu, D. Turner, C.C. Chen, D. Maltz, X. Yang, L. Yuan, M. Zhang. NetPilot: Automating Datacenter Network Failure Mitigation, ACM SIGCOMM 2012 conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM). 2012

[49] R.N. Mysore, A. Pamboris, N. Farrington, N. Huang, P. Miri, S. Radhakrishnan, V. Subramanya, and A. Vahdat. SIGCOMM'09

[50] P. Gill, N. Jain, and N. Nagappan. Understanding Network Failures in Data Centers: Measurement, Analysis, and Implications. ACM SIGCOMM 2011 conference on Applications, Technologies, Architectures, and Protocols for Computer Communication (SIGCOMM). 2011

[51] C. Li, R. Zhou, and T. Li. Enabling Distributed Generation Powered Sustainable High-Performance Data Center. *Int. Symp. on High-Performance Computer Architecture (HPCA)*, 2013

[52] N. Sharma, S. Barker, D. Irwin, and P. Shenoy. Blink: Managing Server Clusters on Intermittent Power. *Int. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, 2011

[53] B. Aksanli, P. Eddie, and R. Tajana. Architecting Efficient Peak Power Shaving Using Batteries in Data Centers. *IEEE International Symposium on Modelling, Analysis & Simulation of Computer and Telecommunication Systems* (MASCOTS), 2013

[54] S.K. Ghai, Z. Charbiwala, S. Mylavarapu, D. P. Seetharam, and R. Kunnath. PC Picogrids: A Case for Local Energy Storage for Uninterrupted Power to DC Appliances, *ACM e-Energy*, 2013.

[55] D. Wang, C. Ren, A. Sivasubramaniam. Virtualizing Power Distribution in Datacenters, *International Symposium on Computer Architecture* (ISCA), 2013

[56] L. Liu, C. Li, H. Sun, Y. Hu, J. Gu, T. Li, J. Xin and N. Zheng. HEB: Deploying and Managing Hybrid Energy Buffers for Improving Datacenter Efficiency and Economy, *International Symposium on Computer Architecture* (ISCA), 2015