Integrated Power Anomaly Defense: Towards Oversubscription-Safe Data Centers

Xiaofeng Hou, *Student Member, IEEE*, Chao Li, *Senior Member, IEEE*, Jinghang Yang, *Student Member, IEEE*, Wenli Zheng, *Member, IEEE*, Xiaoyao Liang, *Member, IEEE*, and Minyi Guo, *Fellow, IEEE*

Abstract— Energy storage devices (e.g., batteries) are critical components for high-availability data center infrastructure today. Without resilient energy management of these devices, existing power-hungry data centers are largely unguarded targets for cyber criminals. Particularly for some of today's scale-out data centers, power infrastructure oversubscription unavoidably taxes the data center's backup energy resources (i.e., UPS), leaving very little room for dealing with power emergency. As a result, an attacker could manipulate the computing system to generate peak power demand and disrupt power-constrained server racks. This study aims at protecting data centers from malicious loads that seek to drain precious energy backup, overload server racks and compromise workload performance. We term such load as Elusive Power Peak (EPP) and demonstrate its basic three-phase attacking model. To defend against EPP, we propose IPAD, a remediation solution build on integrated software and hardware mechanisms. IPAD not only increases the attacking cost considerably by hiding vulnerable server racks from visible power peaks, but also strengthens the last line of defense against hidden power spikes with fine-grained power control strategy. We show that IPAD can effectively raise the bar of power-related attack, with reasonable design overhead.

Index Terms—Data center, power oversubscription, peak power attack, energy storage

1 Introduction

TODAY, data centers are becoming tightly coupled with, and more dependent on energy storage devices (batteries). In the past few years, we have witnessed a considerable interest in deploying *distributed energy backup* (DEB) in data centers from many big companies such as Google [1], Facebook [2], Microsoft [3], and Akamai [4]. This new data center power management paradigm highlights reduced power conversion loss, improved power usage effectiveness (PUE), and much lower total cost of ownership (TCO). According to Akamai, smart batteries placed inside a server or within a rack could drop the required power budget by 14% [5]. In addition, DEB is not only a more energy-efficient alternative to conventional centralized uninterruptible power supply (UPS) system, but also easy to scale and maintain [6]. It could eliminate a potential single point of failure that centralized UPS systems may have.

A DEB-based data center can oversubscribe the power infrastructure without affecting server performance. The basic idea is to shave occasional power demand peaks by discharging a fraction of battery units (no performance capping is performed) [6-8]. Currently there are five ways to deploy batteries in a data center, as shown in Figure 1. The battery unit size varies from hundreds of watts to several MWs. To avoid battery overprovision, normally only one of the five methods is used. Batteries can be installed as top-of-rack UPS, in-server module, and a battery cabinet next to the rack, etc. By directly integrating battery units locally and using DC voltage as backup, one can eliminate energy loss due to double-conversion. Importantly, with DEB a data center can switch (offload) a fraction of server racks to their local energy storage to shave/hide the power peaks at the data center level.

Despite many advantages, power-related security issue could become the Achilles' heel of a DEB-based, aggressively oversubscribed data center. 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 [9-12]. For example, by creating excessive floating-point operations or triggering more cache misses, the attacker can increase system resource consumption considerably [10]. The attacker can also generate simultaneously occurred power surge to overload a system [11, 12].

The security issue turns out to be even worse in many data centers that are heavily power-constrained. If companies continue to squeeze more servers into their existing data center, the risk of power violation may rise rapidly. In addition, DEBs have been frequently used as energy buffer in many green data center designs to handle the power variability [13-15]. In both cases, batteries often experience unusual cyclic usage but do not receive timely recharge. Without enough energy backup, racks are left unguarded from malicious loads. As we transition from centralized to distributed battery architecture, server racks unfortunately become more vulnerable to power anomalies. 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, 16]). 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. By far, the biggest root cause of power outage is battery failure and capacity overload [17], which could have been avoided with a proactive security-aware energy management strategy.

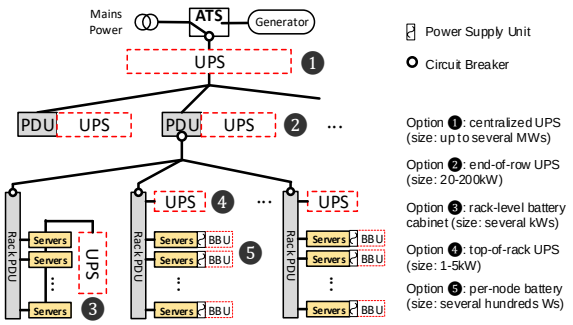

Fig. 1. Major battery deployment methods in a data center

It is important to build a security-aware power management framework since 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 data loss. Unplanned power outage has been shown to cost over $10 per square meter per minute for 40% of the benchmarked data centers [18]. On average, the financial loss of a data center power outage in 2013 is more than $7900 per minute [18]. According to an industry survey, more than 75% data centers require at least 2 hours to investigate and remediate incidents [19]. It means that a successful power-related attack can easily cause the victim data center to lose one million dollars.

In this work we argue that a sophisticated attacker can keep changing its load behavior to create an important type of attack which we refer to as *elusive peak power* (EPP). In the beginning, EPP can create non-offending visible power peaks (disguised as benign loads) to drain the DEB. Afterwards, EPP can be mutated to create various offending hidden power spikes. If there is certain hard power limit on the power delivery path, EPP can leverage very high and narrow power spikes to bring down the victim rack. Whenever there is a soft limit enforced on the load power consumption, EPP can cause unnecessary software performance scaling.

This study aims to demonstrate the vulnerability of battery-dependent data centers and provide an initial yet practical solution. We focus on the question of **how distributed energy backup systems can be gracefully tamed and leveraged to tackle the challenge posed by elusive malicious loads**. This question is very important for data center owners who want to embrace DEB to improve energy/cost efficiency but cannot afford to compromise service availability/performance.

While various technologies are available to protect our servers, the security issue associated with energy/power has been largely overlooked today [12]. It is difficult to defend against power-/energy- related attacks indirectly with existing methods such as intrusion detection and access control. This is because power demand prediction based on load statistics is often resource-consuming and the results are often inaccurate [20]. 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 [19]. Although advanced power metering can be used for real-time analysis, it is not available in many data centers. Fine-grained sampling is also prohibitive since it requires costly implementation of per-server metering. As a result, attackers can often manage to launch power-related attack without prior detection.

In this work we propose integrated power anomaly defense (IPAD), a novel design patch for securing oversubscribed data centers that are backed by DEB. IPAD does not require detailed knowledge of the workload. It is built on lightweight software and hardware mechanisms. It is a performance-aware design that seeks a balance between availability and efficiency.

Specifically, IPAD provides an additional layer of safety in data centers through a novel two-phase power diagnosis and management. In the first phase, IPAD handles the visible peaks through software-based scheduling. Rather than treats each DEB as separate energy backup, IPAD creates a virtual battery pool called *v*DEB 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, IPAD uses a heterogeneous DEB architecture in conjunction with a speculative performance scaling (SPS) strategy to handle the more dangerous hidden spikes. Unlike prior work that mainly focuses on batteries, we leverage a small-scale super-capacitor called μDEB to assist peak power shaving. It can automatically shave most power spikes to avoid circuit breaker tripping. Importantly, the SPS strategy further allows the data center to maintain necessary μDEB power levels. It opportunistically shaves hidden power spikes based on the monitored μDEB conditions to minimize any negative performance impact on servers.

This paper makes the following key contributions:

- We describe a general threat model for power-related attack. We discuss how a sophisticated attacker can leverage elusive power peaks to compromise a DEB-based data center.
- We propose IPAD, an integrated software and hardware design patch based on a three-tier security policy. It can minimize the impact of potential power attacks on vulnerable server racks.

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

2. Background and Motivation

The power delivery/provisioning systems are often the most expensive and time-consuming items in data center design [15]. Data centers may over-provision their servers to achieve the best return on return of investment (ROI). Safely oversubscribing the power infrastructure has become a critical need in data centers due to the very high "power capacity cost" and "power outage cost".

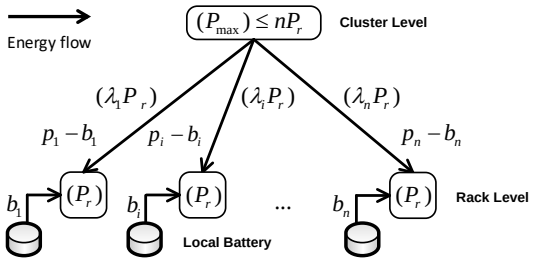

Fig. 2. Power oversubscription model at different power provisioning levels in a scale-out server cluster

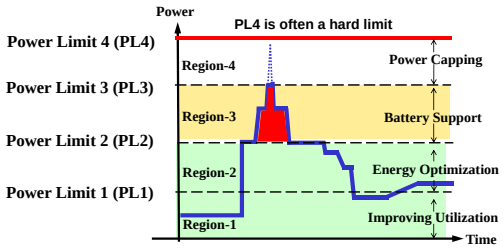

Fig. 3. Threshold-centric power management and the impact of elusive peak power on aggressively provisioned data centers

Table I. Cost, capacity of accuracy, sampling time and application comparison of power metering devices

Power meter name	Power meter type	Accuracy	Time	Cost（per meter）	location
Powerlogic BCPM [22]	Power meters	1%	2-20s	$600-$3,000	Distribution panel
Remote Power Panel [23]	PDU embedded meters	-	-	Included in PDU price	PDUs
Metered Rack PDU [23]	Rack embedded meters	2%-5%	-	$0.04-0.06/watt	Racks
Symmetra PX [28]	UPS embedded meters	-	-	Included in UPS price	UPSs
Masterpact NT [29]	Trip-unit embedded meters	2%	5-60m	$600–$13,000***	Low-voltage switchboards
ION 7650 [30]	Power quality meters	0.02%	1s	$5,000-$11,000***	Utility mains

2.1 Power Oversubscription Model

In this work we focus on aggressively-provisioned data centers (APDC) which leverages power budgeting (capping) strategies to improve utilization. APDC can multiplex the given power budget and enable aggressive server deployment (over-provisioned server resources).

In recent years, energy storage devices (uninterruptible power supply, UPS) are also used for occasional peak power suppression. Without appropriate management, aggressive power provisioning can result in frequent battery usage. Some battery units may incur low levels of stored energy due to uneven discharge. Without timely recharge, we may double the battery usage variation in many cases [12]. These aggressively discharged battery units can be highly vulnerable to power anomalies.

Fig. 2 shows the power oversubscription model of an APDC with a typical two-stage power distribution method. We consider n racks and the power budget of each vertex/edge is given in parentheses. 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 power (nameplate power) demand of each rack is 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 overload, today's intelligent power distribution unit (iPDU) can specify the maximum power of each power outlet. Given the scaling factor$[\lambda_1, \dots, \lambda_n](0 < \lambda_i < 1)$, each power delivery path i can assign a maximum power flow (soft limit) of $\lambda_i P_i$. To avoid overload, we must ensure:

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

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

In general, APDCs today are configured with several power limit (PL) values, as shown in Fig. 3. These power limits logically define four power management regions [21]. For example, in Region-3, power peaks could sustain for limited time, depending on the capacity of local battery. Power Limit 4 (basically the total server nameplate power) is the maximum power demand we can have theoretically. In Region 4, power capping must be used to avoid overload or to extend battery life. Conventional rigid power capping strategies can often cause unnecessary performance scaling on normal tasks (collateral damage triggers by malicious users).

2.2 Peak Power Control Challenges

Existing peak power monitoring and control schemes are not prepared for handling elusive power anomalies. Most data centers lack fine-grained visibility of power variation. The reason is that it requires high cost to deploy plentiful meters or long latency to aggregate power value provided by distributed monitoring programs.

Power Monitoring and Control Cost: Table I summarizes six types of electrical meters mostly widely used in a data center. Power quality meters are expensive devices primarily placed at the utility mains. They monitor crucial electrical parameters such as harmonics and voltage transients. Power meters such as Powerlogic BCPM [22] are less accurate compared to power quality meters. They are often installed on distribution panels to control power circuit loading and balancing. The advantage of power meter embedded in a trip unit is programmable (despite its high cost and low accuracy). Sometimes the UPS and PDU may have built-in meters for computing the PUE. They are generally not verified by national standards and are inaccurate for critical incident alarming.

We expect a range of $50~$230 per kW of IT load depending on how thoroughly the medium-/low- voltage distribution systems are metered [23]. The estimation is conservative since it excludes the cost of electronic trip unit metering, meters embedded in PDUs, etc. Thus, various types of advanced power meters with fine-grained sampling or metering capability is often prohibitive due to the considerable cost to fully sample and meter the upstream electrical distribution in a large-scale facility.

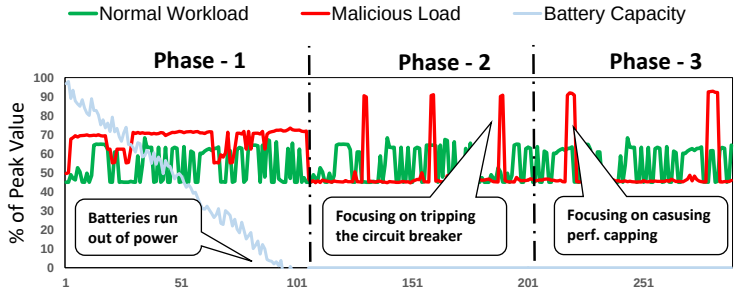

Fig. 4. Demonstration of the three-phase power attack model

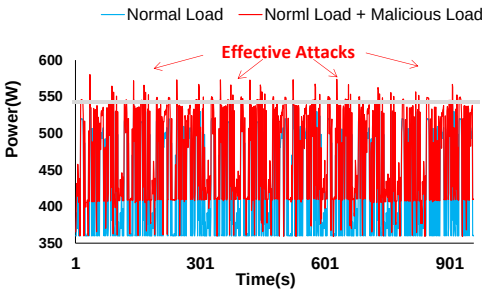

Fig. 5. Demonstration of effective power attack attempts

Fine-grained Power Profiling at Scale: There are several studies on how to acquire per-server power statistics. For instances, Power Containers [24] can perform task-level power supervision and allocation to constrain the overall loads power into the cores' limits. Manufacturers such as Intel [25] and IBM [26] have embedded a set of Machine-Specific Registers (MSRs) into processor architectures and provided various interfaces to monitor the power, energy and thermal status of cores and packages. Taking Linux kernels for example, it can leverage the *perf_event* interface supported by Perf tool to read power-relative data on per-second basis; it can also modify the RAPL interface provided by Intel to obtain power values every 100ms.

Despite the fine-grained monitoring capability, control latency can be a problem due to the complex interactions among nodes during the information gathering process [27]. When power control feedback loops operate at multiple levels in the hierarchy, control stability requires that the lower layer control loop must converge before the upper layer can move on to the next iteration. Therefore, server-level monitoring latency is amplified by scale. This work is partly driven by the need of a more resilient peak power management framework in a non-ideal power monitoring and control environment.

3. Threat Model

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

We propose a three-phase attack tailored to the power behaviors of today's oversubscribed data centers. Specifically, the attack is organized as three steps. The first step is not offensive. The second and the third steps are offensives acts that have different goals. With the above elaborated efforts, a sophisticated adversary can manipulate the power demand of a small group of compute nodes to overload/overwhelm a larger server cluster. We call this act as *Elusive Power Peak* attack.

To overload the server rack and trigger circuit breaker the attacker first needs to subscribe a few physical servers. These machines will become the hosts of the malicious loads. The attacker can opportunistically look for such a host by repeatedly creating many virtual machines (VM) and monitoring the IP of the VM instance. One can also keep rebooting a few VMs until they research the same desired location [24]. Once the attacker has gained control of enough nodes, the next thing is to wait for the best time to change its power demand.

3.1 Phase-I Attack: Disguised Power Peak

Servers with inadequate stored energy are much easier to overload. Therefore, the attacker first needs to create or identify vulnerable racks by initiating a "*disguised power peak*" which can mildly increase the average utilization of the server rack. In most cases, the data center will treat such power demand as normal load surge. This phase represents the latent period of the power attack.

Figure 4 demonstrates this process using realistic battery-backed server clusters. In Phase-I, the attacker keeps running workload to accelerate battery discharge. These local batteries become temporarily unavailable since most DEBs choose to disconnect low-power batteries from load for safety reasons. For example, Facebook uses a disconnect device to isolate battery if the sensed terminal voltage drops below 1.75V per cell [2].

The attacker become aware of battery status by monitoring workload performances. Once the peak-shaving DEB runs out of power, one must use power capping techniques such as dynamic frequency scaling (DFS) for emerging handling purpose. By monitoring the execution latency or similar metrics the attacker could be able to identify when and where the stored energy become unavailable. 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 the "disguised power peak".

3.2 Phase-II Attack: Offensive Power Spikes

In this stage, the attacker can start to launch "*offensive power spikes*" that will create power surges possibly invisible to data centers. In other words, once the attacker drains the batteries with invisible peaks in the Phase-I, EPP has been upgraded to an aggressive power overload attack. Otherwise, these local batteries can eliminate any power surge including fine-grained spikes.

As shown in Figure 4, the peak power virus can be mutated to create very high and narrow power spikes in Phase-II. The power spikes are considered "hidden" since they are short load surges which do not significantly increase the average utilization. Existing utilization-based power monitoring mechanisms cannot detect such fine-grained variation [20]. They normally monitor the total energy consumption at coarse-grained intervals (e.g., 10

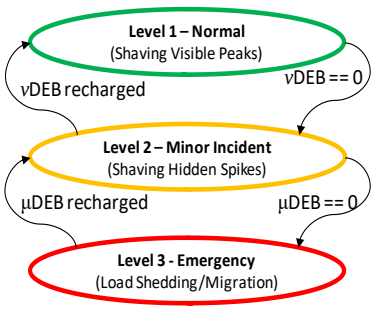

vDEB >0?	μDEB >0?	VP >0?	Initial State
0	0	0	L3
0	0	1	L3
0	1	0	L2
0	1	1	L3
1	0	0	(L1/L2)
1	0	1	(L1/L2)
1	1	0	L1
1	1	1	L1

Fig. 6. 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 system

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. Advanced power capping can operate much faster, but it mainly works on per-node level. In fact, within a power oversubscribed environment, each server is allowed to reach its peak power as long as the total rack/PDU utilization is within the budget.

Whether or not an effective attack can trip the circuit breaker depends on the actual over-current and the peak current duration [20]. Tripping a circuit breaker is not an instantaneous event since most PDUs 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. 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. Given enough overload events, it has very good chances to fail a server rack.

3.3 Phase-III Attack: Offensive Power Peaks

Note that the "offending hidden power spikes" will not guarantee a successful attack (i.e., power failure). The way the attacker launches power spikes greatly affects attacking results. We consider three key factors: the height, width, and frequency of power spikes. As the attacker uses more aggressive attack approaches (increasing spike duration, frequency, etc.), the chance of being detected by the data center also increases. An effective attack does not equal to a successful attack; it means that the load power exceeds the provided power budget. It only increases the probability of power outage. As discussed in our prior work [12], the attacker may create effective attack even when the power budget is adequate (e.g., 96% of the nameplate power).

If the attacker cannot trigger overload in the Phase-II, it can be degraded to a mild power abuse attack. In this phase, the attacker does not have to subscribe a large number of physical servers. One can only launch DDoS-like attacks with power-hungry queries. This is very similar to the power grab concept [21] in power-constrained data centers. In this case, the attackers present existing power management framework with an embarrassing situation. If we ignore it and use conventional rigid power capping strategy, peak power shaving activities can at the same time cause collateral damage (unnecessary performance scaling) on normal tasks. This includes slowdowns of normal user tasks and frequent depletion of the newly re-charged precious backup energy. Both can disrupt normal operation and render the cost-saving effort of current power management framework ineffective.

4. Integrated Defense Mechanism

To tackle the security challenge faced by existing oversubscribed data centers, we propose *integrated power attack defense* (IPAD). IPAD is a new design patch to exiting power management framework. It allows data centers to run safely and smoothly under power anomalies.

4.1 Basic Policy

IPAD defines a three-tier security policy for power management on battery-backed server clusters. It lays down the general rule for protecting data centers from malicious loads that intend to overload the system.

IPAD 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, where VP indicates if a visible peak is identified.

As shown in Figure 6, our policy defines the states for all the combinations of initial inputs. Depending on the operating environment, IPAD may enforce different security levels and expose underlying power/energy profile to the data center. This allows one 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 IPAD 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 *v*DEB and μDEB are empty, IPAD will overlook the load power behavior and force to enter an emergency state. This can cause the data center to lower server performance or shed loads, e.g., put some servers into sleeping/hibernating states. Although the temporary load shedding may incur certain performance degradation, it is not overkill. This prevents data center from incurring significant loss during a large-scale power failure. In fact, by sleeping only a small number of servers, one can prevent most of the 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.

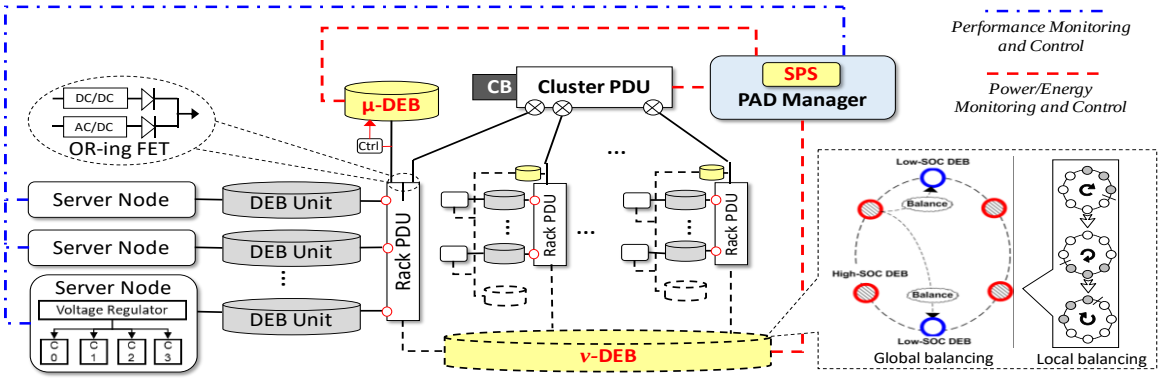

Fig. 7. The IPAD architecture. In the figure it shows the deployment of three critical components, i.e., vDEB, uDEB, and SPS module

4.2 IPAD Design

The main source of vulnerability lies in the use of traditionally simple, homogeneous energy storage architecture to defend against a potentially sophisticated peak power anomaly. In the following we first discuss the IPAD's hybrid energy backup design for prolonging the survival time of data centers. Afterwards, we discuss speculative performance scaling, an optimization strategy that can lower the chance of entering the Level-3 state.

4.2.1 Hybrid Distributed Energy Backup

Our IPAD design aims to manage complex power anomalies. As shown in Figure 9, it exploits the energy stored in both batteries and super-capacitors. The batteries form a virtualized energy backup pool called virtual DEB (vDEB). The super-capacitors form a micro energy backup pool called micro-DEB (μDEB). The vDEB module aims at protecting data centers from a brute visible peak attack in the Level-1 emergency state. The μDEB design intends to defend against a more sophisticated hidden spike often seen in the Level-2 state.

Virtual DEB

Rather than treats rack-mounted batteries as separated energy backup systems, IPAD creates a virtual energy backup pool termed vDEB. During the runtime, our vDEB management strategy allows server racks to share unused energy backup capacity within the same PDU.

In conventional designs, some racks may aggressively discharge their batteries and at certain point they 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 operation of servers can be disrupted.

Our vDEB energy usage balancing strategy combines cluster-level battery balancing and rack-level battery balancing. As shown in Fig. 7, we assign the discharge rate of each battery unit based on its available SOC value rather than the loading conditions of racks or PDUs. It prevents vulnerable batteries from aggressively discharging and allows for fast balancing. Implement vDEB does not require us to rewire the data centers since current racks or PDUs generally support power switching between batteries and utility supply. Besides, we also set an upper bound when assigning the discharge rate (i.e. represented by the ideal discharge power P_{ideal}). It could prevent accelerating the aging process of battery systems.

The vDEB design brings two 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 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 of battery usage).

Micro DEB

Virtual DEB alone cannot defeat a well-planned power attack. A peak power virus can be mutated to create transient power spikes that most utilization-based power management software can hardly detect.

PAD further integrates a dedicated small energy backup device to existing distributed battery system (Figure 7). The device, termed as micro DEB (μDEB) in this work, is designed to further strengthen the defense against hidden power spikes at the server rack level. To shave the hidden power spikes, the μDEB must be able to react to any voltage surge/sags. We connect μDEB with the primary power delivery bus using an ORing controller, 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 power supply unit. Second, at the rack level, long time current sharing for sustained peak shaving can cause thermal issues.

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 en-

ergy capacity. Normally a 2A battery cell can provide 85 W at the maximum for 5minutes [30]. This requires us to connect many battery cells in parallel to achieve the desired power capacity, which can be bulky and expensive.

4.2.3 Speculative Performance Scaling

From the above discuss we can see that IPAD creates a virtual energy backup pool to increase attacking cost and introduces a dedicated energy backup device to handle undetectable spikes. However, even μDEB cannot prevent a data center from entering the power emergency state (see Fig. 9, the Level-3 emergency state). It only extends the survival time. The stored energy in a μDEB decreases as a sophisticated attacker keeps generating hidden spikes. Once the μDEB runs out of power, one can hardly detect the hidden peak power virus any more. If one keeps using aggressive power capping to maintain the stored power level of μDEB, it can cause significant performance degradation. Thus, it is crucial to carefully utilize μDEB and smartly discharge the energy in it.

A more notable feature of IPAD is that it can adaptively adjust its peak power capping strategies in Phase-III for better design tradeoffs. In general, if the remaining μDEB capacity is adequate, PAD allows the system to run at full speed for performance considerations. During runtime, it speculates about the position of the power spikes. Upon insufficient uDEB capacity, PAD throttles CPU performance intermittently to limit the server power.

Given a continuous set of peak power viruses A_i, PAD coordinates load behavior and SC behavior to handle different situations. We use $A_i = \{aw_i, ah_i, af_i\}$ to represent the peak power virus which triggers the i^{th} μDEB discharge. Here, aw_i, ah_i and af_i are respectively the weight, height and frequency of the attack. The riskiest malicious vector for an attacker is $A_{max} = \{A_W, A_H, A_F\}$. Namely, the ability of the attacker will never exceed A_{max}, otherwise, it will expose itself to the system power managing center.

Our design does not aggressively throttle CPU frequency. Instead, we propose a *speculative performance scaling* (SPS) method. SPS speculates the possible frequency of power spikes and it intends to perform DVFS only when the power spike arrives. It allows us to achieve a better tradeoff between security and performance.

SPS considers several important questions such as when to adjust server power states, how much to lower, and how long to sustain the scaled speed. We configure PAD with a set of metrics $SPS_i = \{start_i, level_i, stop_i\}$. Here, $start_i$ and $stop_i$ means the start time and end time of the next DVFS event, i.e., the duration of performance scaling. $Level_i$ means the designated performance level. At the i^{th} discharge events of μDEBs, PAD updates its DVFS configuration with $SPS_i = \{start_i, level_i, stop_i\}$.

In Figure 9 we show different scenarios of SPS. At the beginning, server nodes run at full speed, i.e., $SPS_0 = \{0,100\%,0\}$. Correspondingly, the initial μDEB energy level is E_{max}, capable of covering the most extreme hidden surges, i.e., A_{max}. Once the remaining μDEB energy becomes lower than a predefined threshold E_{low}, SPS is obliged to handle the invisible power spikes by confining overall consumed power to the safety line. After μDEB has over E_{high} electricity, SPS reconfigures the system to ensure high performance.

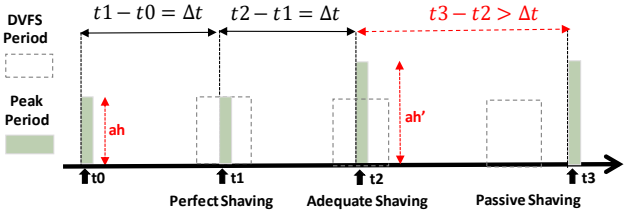

Fig. 8. Illustrative example of SPS

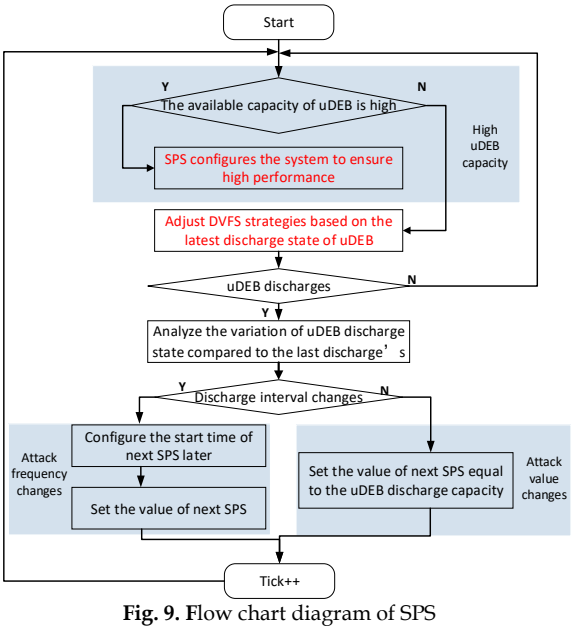

Fig. 9. Flow chart diagram of SPS

With speculative performance scaling, IPAD examines its DVFS strategy whenever μDEB discharges. Namely, the reconfiguration of DVFS is determined by the change of μDEB energy level that indicates the existence of hidden power spikes. In this scenario, the next DVFS configuration is calculated during the waiting period based on the last monitored μDEB energy footprint.

In Figure 8, when the attacker arrives at time t_0, μDEB discharges to eliminate the spikes. DVFS doesn't work because SPS_0 is {0,1,0}. Thus, attacker A_0 considers itself unknown to the system detection center since there is no change of the system's execution speed. Thus, it's more likely for the attacker to strengthen its attack power or keep the same attack method for the next time.

The discharge of μDEB will cause SPS to enter into a more aggressive mode. If EPP keeps the same at t_1, the pre-configured system exactly eliminates it (the power peaks and DVFS period overlaps). Since the μDEB does not discharge, there is no need to reconfigure the DFVS as well. Meanwhile, the attacker can increase its peak power height as shown at t_2. The proposed SPS can only eliminate partial power peaks in this case. IPAD must rely on μDEB to eliminate the rest attack peaks which result in μDEB discharge. Once the attacker knows that it has been identified according to the observed voltage/frequency changes, it can decrease its attack frequency and trigger a spike at t_3. In this case, the SPS may temporarily lost its

Table II. Evaluated Applications

Workload Name		Workload Type	Target	Benchmark	Description
W1	*John the Ripper*	Cryptography Security	Processor	PST	Password cracker programs.
W2	*Encode MP3*	Audio Encoding	Processor	PST	MP3 encoder licensed under the LGPL.
W3	*Ray tracing*	3D Graphics	Processor	BDB	Creating 3D graphics using ray-tracing.
W4	*TSCP*	Games	Processor	PST	Tom Kerrigan's Simple Chess Program
W5	*Sand*	Bioinformatics	Processor	BDB	Accelerating genome assembly.
W6	*IOZone*	IO Programs	Disk	PST	Testing file system, disk performance.
W7	*Postmark*	NetApp's PostMark	Disk	PST	Simulating web and mail service.
W8	*Stream*	RAM test program	Memory	PST	Testing the RAM performance.
W9	*Loopback*	TCP micro-benchmarks	Network	PST	Testing network adapter performance.
W10	*Recommendation*	E-commerce	System	BDB	Predicting the preferences of the consumers

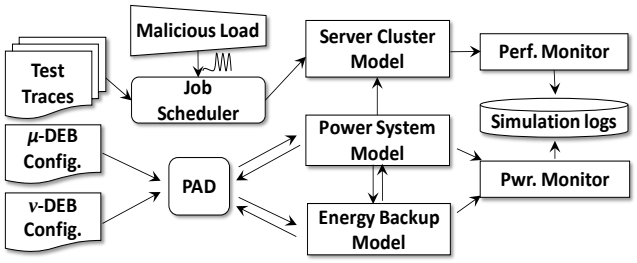

Fig. 10. Trace-based validation framework. We feed realistic workload traces, data center power system data, and power anomaly data into our simulators

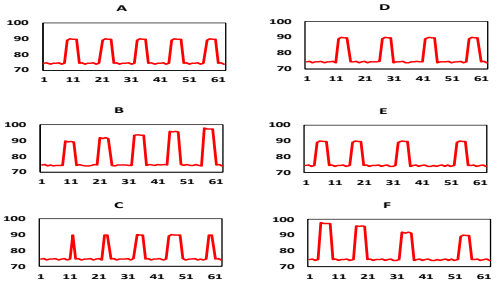

Fig. 11. Example of the collected attacking traces. A: dense and constant. B: height-varying. C: width-varying. D: sparse and constant. E: frequency-varying. F: height, width, frequency-varying attack

Table III. Evaluated power management schemes

Scheme (Abbr.)	Description
Load Shedding (LS)	Conventional techniques that only use DVFS mechanisms to cap power peaks
Peak Shaving (PS)	Emerging designs which leverage distributed batteries to shave power shortage
Basic Defense (PAD)	Uses vDEB to extend battery life and uses μDEB to shave hidden power spikes
Advanced Defense (IPAD)	Optimized PAD that combines PAD with speculative performance scaling

peak power shaving capability due to misprediction.

In general, IPAD can always mitigate the impact of hidden spikes if they demonstrate certain patterns (as Figure 11 shows). Even if SPS is not enough to handle the power spike it can still weaken the attack because SPS makes the attack believe that it has been identified. Once the adverse nodes change attack frequency, IPAD will not be able to precisely track the power spike. In this situation, IPAD must adaptively re-compute DVFS configuration according to μDEB. Note that autonomous power capping at the server-level can further aid IPAD to defend against aggressive power attack. At this moment, our proposed IPAD design does not consider per-server autonomous power capping since it can compromise the overall server utilization without a global view of power usage planning. Meanwhile, coordinating per-server power manager can be complicated with large number of nodes in a data center.

5. Evaluation Methodology

We build a scaled-down testing platform as shown in Figure 13-A. It consists of 2 mini server racks with 10 server nodes and a set of three YUASA UPS batteries per rack. The power capacity of batteries per rack is 1000W to ensure that it can maintain 10 minutes under full load. Each rack also contains several Maxwell BCAP0650 P270 Ultra-capacitors with maximum energy capacity of 11250 Joules. All the batteries and super-capacitors are dynamically monitored on a per second basis. Our system can dynamically switch ON/OFF the UPS with SNMP commands over Ethernet and collect key battery and power information during runtime. The UPS batteries are discharged whenever a SNMP command asking them to switch into discharging state because of the power demand exceeding the target budget. But the batteries only recharge after every complete discharge considering its limited cycle lifetime. Ultra-capacitors discharge whenever the power demand exceeds the budget without any operation commands. SCs can also charge whenever the power demand is lower than the budget, because SCs are not constrained by the charge/discharge cycling.

We model different power viruses taking advantage of stresscpu benchmark under Phonorix Test Suites. 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 14 shows power virus trace examples we generated. Based on the configuration of our system, we consider both invariant and variant power attacks. The invariant virus trace keeps the same power attributes (e.g., width, height) discussed in Section 3 through its whole attack process. Instead, variant attack will continuously assess the data center power detection and management schemes. It then intelligently tunes attack vectors in accordance to its exposure rate. We also consider two types of power attacks: dense and extensive

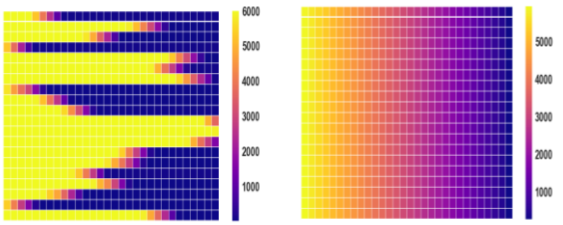

Fig. 12. A comparison of DEB usage in conventional datacenters (left) and IPAD (right). x-axis: seconds; y-axis: rack ID

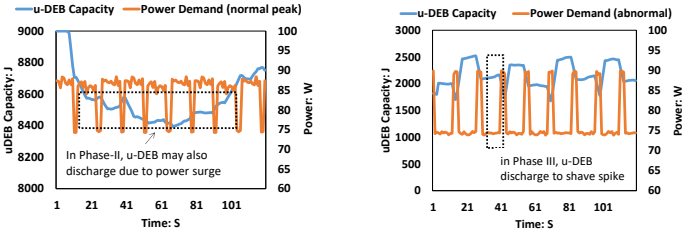

Fig. 13. *u*DEB discharging states. Left: discharging mainly for eliminating malicious spikes. Right: discharging for supporting other power schemes

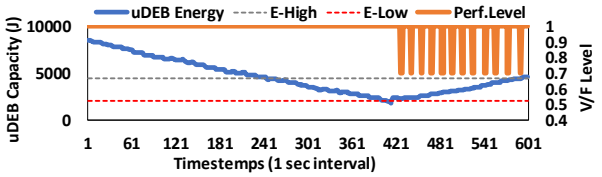

(A) Fixed attack size or arrival time

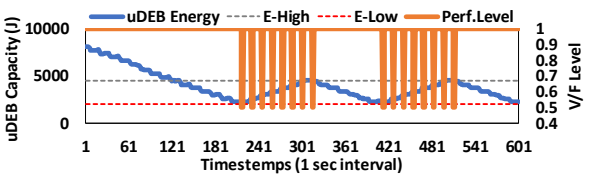

(B) Variable attack size

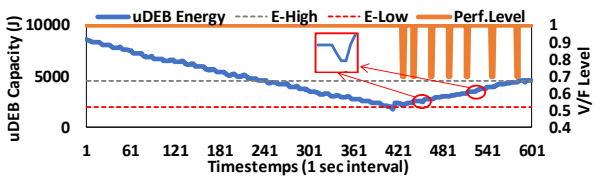

(C) Variable attack arrival time

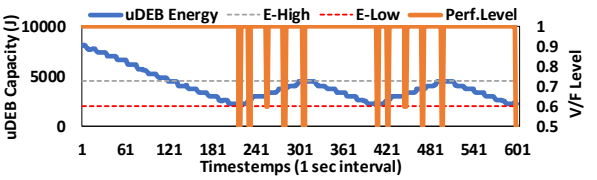

(D) Changing attack size and arrival time

Fig. 14. The *u*DEB energy usage pattern of IPAD under different attack scenarios. The solid blue line represents *u*DEB energy and the solid orange line means server performance.

power spikes as well as sparse and less aggressive spikes.

We feed the collected power virus traces to a trace-based data center simulator. At the same time, this simulator takes other compute traces as input. These traces are collected through running the most emerging applications (as shown in Table III) in present largescale systems. We choose ten mainstream applications from BigDataBench [32] and PST [33]. In our simulation platform, we assume that each server has five discrete frequency/voltage (F/V) scaling levels: 1GHz/1V, 0.9GHz/0.9V, 0.8GHz/0.8V, 0.7GHz/0.7V, 0.6GHz/0.6V. We select the F/V settings according to the system's loading state.

We verify whether IPAD can prevent a more intelligent attacker. Namely, the attacker may adjust its attack trace in phase-II. It can always keep same (as shown in Fig.10-(a)~(b)) or changes its attack force (as shown in Fig.10-(c)~(d)) and attack frequency (as shown in Fig.10-(e)~(f)) according to the system's reaction. We explore the cooperation modes between μDEB and SPS when handling these attack scenarios.

6. Results

In this section we evaluate the impact of IPAD on power-constrained data centers. Table 3 summarizes the power management schemes we have evaluated.

6.3 DEB Utilization

We first examine the behavior of energy storage devices under IPAD. We must ensure PAD doesn't compromise their normal functionality. Since the main function of μDEB is to detect and handle invisible peaks, we also examine whether the visible peaks can waste them.

Fig. 12 shows the monitored DEB utilization map of the evaluated server clusters at each fine-grained timestamp. In the figure, light yellow represents fully charged batteries while dark blue means near-empty batteries. Racks with low energy storage backup could be ideal targets for a sophisticated criminal. With conventional design, some server racks 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. In contrast, *v*DEB allows a data center to hide vulnerable racks by balancing battery usage. Although uneven usage still exists to some extent, those server racks no longer differ significantly in their backup power at any timestamp.

To verify that μDEB mainly takes charge of detecting and handling invisible peaks, we design two scenarios. The first scenario is an attack scenario in which the attacker generates invisible peaks. As shown in Fig.13-(a), *u*DEB can handle the invisible peaks and then get recharged. In Fig.13-(b), there are no attack nodes, but the good nodes generate some high peaks. These peaks will drain all the energy stored in the *u*DEB in a short time (the *u*DEB capacity decreases), In this case, the system will deal with them using mature power shaving strategies such as DVFS and processor power gating.

6.2 Control Effectiveness

It is important for IPAD to be able to handle power outage even under the extreme situation, i.e., Level 3 with variable power spikes. IPAD intends to smartly trigger

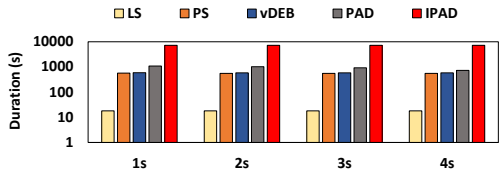

(A) Different peak power widths

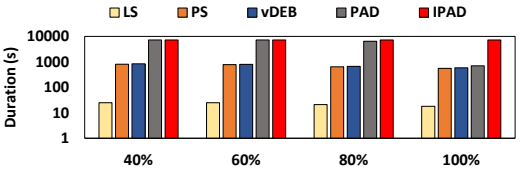

(B) Different ratios of victim servers

Fig. 15. The sustained operation duration under various power attacks

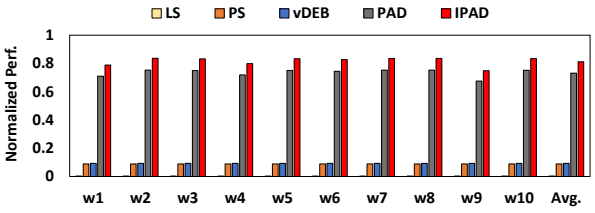

Fig. 16. Workloads' performance under different schemes

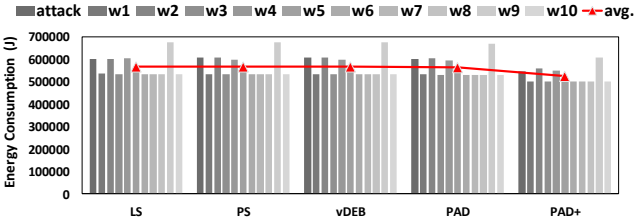

Fig. 17. The evaluated workloads energy consumption

performance scaling (via DVFS) to smooth power demand. It dynamically adjusts its power capping strategy based on the monitored μDEB capacity.

Fig.14-(a) shows the results under the simplest attack scenario: an attacker maintains regular power spikes. The attack width is 3 seconds and its frequency is once every 14 seconds. It is easy for our SPS to handle all the invisible peaks even if μDEB is insuf ficient. As we can see, IPAD allows data center to perform intermittent power capping. Compared to exiting techniques, IPAD can lower the performance impact of power capping.

Adjusting power attack height and power attack timing has different impacts on data centers. Increasing the attack height can trigger deeper performance scaling. In addition, μDEB energy still runs out faster in Fig.12-(b). Nevertheless, the μDEB is charged more quickly due to the more aggressive SPS configuration. Adjusting the attack intervals may degrade the attacking effectiveness. It is obvious that the attack spikes get sparser, as shown in Fig.12-(c). If the attacker changes its attack initiate interval (it changes the frequency at the interval of 10 seconds), it has little chance to escape SPS since IPAD is be able to track the position of the power spikes.

However, when combining the two approaches, it may have devastating effect on data centers. Fig.12-(d) shows the result when the attacker changes both the launch time and value. The malicious trace eventually exhausts μDEB capacity and causes significant performance degradation through the maximum attack vector.

6.1 Load Survival Time

We evaluate the time for which the critical loads may sustain during a power attack. We define survival time as the duration between the start of the attack and the first overload event. Figure 15 presents five different power management schemes under different attack scenarios.

Although the attacker may change its attack attributes, IPAD is more likely to defeat the attack. In other words, IPAD effectively prevents power outage at the early stage of power attack. PAD can increase the sustaining time, but power failure still happens as the μDEB devices run out of power. Existing schemes cannot stop the phage-II attack when the attack force is getting stronger.

Figure 15 also indicates that μDEB and vDEB have different impact on data center survival time. Compared to μDEB, the improvements of vDEB are bigger. This is mainly because that the visible power peaks dominate in the overall attacking period. Combing μDEB and vDEB allows PAD to better deal with various power virus. Overall, PAD improves the sustained time by an order of magnitude compared to conventional data centers.

We recognize that PAD cannot 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. We evaluate a small cluster and therefore the results are not striking. In data centers that have hundreds/thousands of racks, IPAD can offer impressive security/availability benefits.

6.4 Performance and Efficiency

It is important that our security-conscious power management framework does not greatly compromise the performance of normal tasks. Since IPAD only triggers capping when attack happens, it has little side-effects on workload performance during normal operation.

We monitor the degree of server voltage and frequency scaling to quantify the impact of various power management schemes on task performance. In Fig 16, the Y-axis represents normalized performance and the X-axis shows the workload described in table III. It is beneficial to deploy speculative performance scaling (SPS) other than very strict power capping. We evaluate our design under given power attack in Phases II and III. As we can see, IPAD provides the best performance guarantee compared to the other schemes. Conventional power management mechanisms such as LS and PS cannot tackle Phase-II attack; thus, their performance is very low. Differently, since PAD and IPAD both use μDEB and vDEB to manage peak power, they are orthogonal to existing system and software level power optimizations. In addition, the speculative server load scaling approaches can reduce

unnecessary power capping activities that are seen in other baseline. Overall, PAD could yield much improved performance (by 8X) and IPAD further boosts the performance by 11% on top of it.

Finally, it is also attractive for a security-oriented design to be energy-efficient. We compute the total consumed energy with the consideration that it is sourced from batteries, ultra-capacitors, as well as utility power supply. In this experiment, we run different power management schemes for 2 hours with a total power budget of 90% of the nameplated power. The peak power virus consists of 50% of the total server nodes and the attack width is 4 seconds. For schemes without μDEB, power failure happens once the Phase-II attack occurs. Thus, we only compute their energy consumption in Phase-I. For the other schemes, we calculate the aggregated energy consumption in the whole observation period.

As shown in Figure 17, PAD and IPAD consume less energy since they leverage adaptive power throttling to decrease the server nodes' power if Phase-II attack virus exists. Compared to PAD, IPAD is slightly more efficient (7% less energy) as it leverages adaptive performance scaling to decrease μDEB consumption. It means less operational cost (OpEx) or the possibility of adding more machines into existing data center infrastructure.

6.6 Sensitivity Analysis

Increasing the capacity of uDEB allows one to better shave the hidden power spikes, and therefore increasing survival time. However, since the SPS algorithm can mispredict power spikes, uDEB still undergoes energy loss during aggressive attack. As shown in Figure 18, if we increase the capacity of uDEB, the survival time increases linearly. Differently, the normalized performance is not proportional to the capacity growth of uDEB. Increasing uDEB improves performance greatly in the beginning. As uDEB becomes larger, the workload performance eventually becomes relatively stable. This is because SPS changes uDEB from speculative capping and aggressive capping. When stored energy of uDEB is low, the percentage of aggressive capping is high, which greatly increases performance overhead. When the uDEB is full of power, the system does not frequently trigger power capping.

Note that larger uDEB is not an ideal design choice. 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. We do not treat vDEB as cost overhead since IPAD leverages battery devices that most data centers already have. 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 IPAD. We expect that data centers may adopt different capacity planning strategies to achieve their desired trade-offs in profitability, availability, and security.

7. Related Work

This section discusses prior studies in different domains that are most relevant to our work.

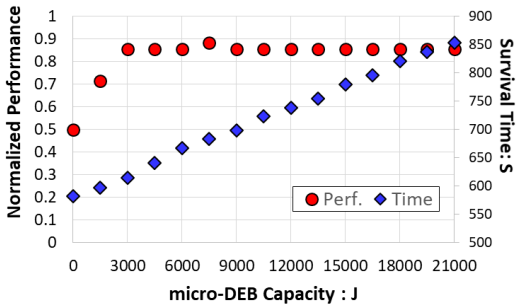

Fig. 18. Impact of uDEB capacity on performance and survival time

7.1 Power/Energy Related Attack

Power and energy related attacks are drawing increasing attentions. Vulnerability in server power management framework has been identified recently in terms of energy abuse attack [9, 10, 34, 35] and power overload attack [11, 12, 36, 49]. An energy abuse-based attack mainly intends to consume additional server energy. Differently, a power overload-based attack aims to cause rare but expensive outages. Islam et al. [36, 37] demonstrate that an attacker can launch well-timed power attacks with a high successful rate to trigger outage but does not consider the energy backup. In contrast, we look at DEB system in the data center. We investigate a new attack approach which intends to overload servers that have limited backup.

7.2 Resource Availability Attack

Another representative group of related work is in the context of resource contention, especially in consolidated data centers [38]. Several papers have investigated the availability issues with regard to computing/networking systems [39, 40, 41, 42]. At the chip level, hardware Trojans can be used to block a network-on-chip system, causing denial of service [41]. At the server system level, a resource-freeing attack (RFA) could modify a victim VM's workload [42]. According to prior research, attackers can affect the availability of power resource at the facility level through running intensive loads [11, 12, 35] on the controlled VMs or invoking frequent VM mitigation activities [21]. As power budget shrinks, Power Grab [21] can abuse power resources and disrupt operation of its competitors by operating power-hungry VMs. Different from these works, we focus on running task-intensive VMs can drain the precious energy storage and overload server racks without prior detection.

7.3 Aggressively Provisioned Data Center

Designing aggressively provisioned data center has attracted great attention in the past several years [6, 7, 12, 43, 46]. Since servers rarely reach peak load simultaneously, power over-subscription shows great promise in maintraining data center performance scaling trend with attractive cost efficiency. To ensure that the power dissipation stays below the given power budget, aggressive control strategies such as power/performance state tuning [27, 44, 45] and battery-based peak power shaving [7, 12, 43] are employed. Performance-preserving aggressive

power capping framework has been deployed in the industry [47]. However, current power management frameworks are too optimistic. They merely consider utilization and performance in the normal scenarios with overlooking malicious power activities. We illustrate that a sophisticated attacker can exploit the blindness of these power management schemes to mount an attack.

7.4 Battery Management

Conventionally, batteries are only used as emergency backups which are always centrally deployed and rarely used; or they 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 [12, 13, 14, 48, 49, 50]. Some recent proposals have focused on managing distributed batteries [6, 32, 42, 51, 52, 53], which can be used to shave the occasional load power peaks [7, 8] in aggressively provisioned data centers. However, they only focus on energy/power efficiency and does not consider the security issue of aggressive power management. As a result, the associated servers are often the potential victims of power virus.

8. 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 new power management scheme 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 with better performance guarantee and negligible cost overhead.

Acknowledgement

This work is supported, in part, by the National Basic Research Program of China (973 Program, No. 2015CB352403) and the National Natural Science Foundation of China (No. 61502302, 61702329). Corresponding author is Chao Li from Shanghai Jiao Tong University.

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] Akamai: How batteries could cut datacenter power bills, https://www.zdnet.com/article/akamai-how-batteries-could-cut-datacenter-power-bills/

[5] 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.

[6] 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.

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

[8] 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.

[9] F. Palmieri, S. Ricciardi, U. Fiore, M. Ficco, and A. Castiglione, "Energy-Oriented Denial of Service Attacks: An Emerging Menace for Large Cloud Infrastructures", *The Journal of Supercomputing* (TJSC), Volume 71, Issue 5, pp 1620–1641.

[10] Z. Wu, M. Xie, and H. Wang, "On Energy Security of Serer System", *IEEE Transactions on Dependable and Secure Computing* (TDSC), Volume 9, No. 6, 2012.

[11] 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.

[12] C. Li, Z. Wang, X.F. Hou, H. Chen, X. Liang, and M. Guo, "Power Attack Defense: Securing Battery-Backed Data Centers", *International Symposium on Computer Architectur* (ISCA), 2016.

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

[14] 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.

[15] 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.

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

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

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

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

[20] 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.

[21] X. Hou, L. Hao, C. Li, Q. Chen, W. Zheng, and M. Guo. "Power Grab in Aggressively Provisioned Data Centers: What is the Risk and What Can Be Done About It". *Proc. the 36th IEEE International Conference on Computer Design* (ICCD), Oct. 2018.

[22] PowerLogic power-monitoring units Technical datasheet Branch Circuit Power Meter (BCPM). http://www2.schneider-electric.com/resources/sites/SCHNEIDER_ELECTRIC/content/ live/FAQS/239000/FA239501/en_US/BCPM_Installation_Guide_Z205396-0H.pdf.

[23] D. Kidd and W. Torell. Types of Electrical Meters in Data Centers. Whitepaper 172. http://www.p3-inc.com/images/electrical_meters_data_centers.pdf.

[24] K. Shen, A. Shriraman, S. Dwarkadas, et al, "Power containers: an OS facility for fine-grained power and energy management on multicore servers[J]", *Architectural Support for Programming Languages and Operating Systems* (ASPLOS), 2013, 48(4): 65-76.

[25] Intel® 64 and IA-32 Architectures Software Developer's Manual, Volumn 3B, System Programming Guide, Part 2. Intel processor datasheet. https://www.intel.com/content/dam/www/public/us/en/documents/manuals/64-ia-32-architectures-software-developer-vol-3b-part-2-manual.pdf.

[26] POWER8 Processor Datasheet for the Single-Chip Module. https://www.setphaserstostun.org/power8/POWER8_ds_dd2.x_v17_05APR2016_pub.pdf.

[27] A. Bhattacharya, D. Culler, A. Kansal, S. Govindan, and S. Sankar, "The need for speed and stability in data center power capping", *In Proceedings of the Green Computing Conference (IGCC)*, 2012.

[28] Symmetra PX 250/500//Scalable from 100 kVA to 500 kW, parallel capable up to 2,000 kW. Schneider-electric technical report. http://www.apc.com/salestools/BMOE-7PDSSB/BMOE-7PDSSB_R5_EN.pdf.

[29] Masterpact NT and NW LV power circuit breakers and switch-disconnectors. https://www.schneider- electric.co.kr/documents/ Catalogue/MasterpactNTNW_catalogue.pdf?c_type=emailsend.

[30] ION7550/ION7650 Functions and characteristics. http://www2.schneider-electric.com/resources/sites/SCHNEIDER_ELECTRIC/content /live/FAQS/327000/FA327825/en_US/PLSED306011EN%20PRINT.pdf.

[31] 12v 12Ah Lead Acid Battery, http://www.micropik.com/PDF/CP12120.pdf

[32] L. Wang, J. Zhan, C. Luo, Y. Zhu, Q. Yang, Y. He, W. Gao, Z. Jia, Y. Shi, S. Zhang, C. Zheng, G. Lu, K. Zhan, X. Li, and B. Qiu, "BigDataBench: A big data benchmark suite from internet services", *IEEE 20th International Symposium on High Performance Computer Architecture* (HPCA), Page: 488-499, 2014.

[33] Phoronix Test Suit, http://www.phoronix-test-suite.com/?k=home.

[34] Z. Wu, M. Xie, and H. Wang, "Energy Attack on Server Systems", *USENIX Workshop on Offensive Technologies* (WOOT), 2011.

[35] F. Palmieri, M. Ficco, and A. Castiglione, "Adaptive Stealth Energy-Related DoS Attacks against Cloud Data Centers", *Eighth International Conference on Innovative Mobile and Internet Services in Ubiquitous Computing*, 2014.

[36] M. Islam, L. Yang, K. Ranganath, and S. Ren, "Why Some Like It Loud: Timing Power Attacks in Multi-tenant Data Centers Using an Acoustic Side Channel." *International Conference on Measurement and Modeling of Computer Systems* (SIGMETRICS), 2018.

[37] M. Islam and S. Ren, "Ohm's Law in Data Centers: A Voltage Side Channel for Timing Power Attacks", ACM *Conference on Computer and Communications Security* (CCS), 2018.

[38] J. Liu, S. Wang, A. Zhou, J. Xu, and F. Yang, "SLA-Driven Container Consolidation with Usage Prediction for Green Cloud Computing", *Frontiers of Computer Science* (FCS), Vol 14, Issue 1, 2020

[39] D. Grunwald, and S. Ghiasi. "Microarchitectural denial of service: insuring microarchitectural fairness", *International Symposium on Microarchitecture* (MICRO), 2002.

[40] D.H. Woo, and H.S. Lee, "Analyzing Performance Vulnerability due to Resource Denial-of-Service Attack on Chip Multiprocessors", In Workshop on Chip Multiprocessor Memory Systems and Interconnects, 2007.

[41] T.H. Boraten, and A.K. Kodi, "Mitigation of Denial of Service Attack with Hardware Trojans in NoC Architectures", *IEEE International Parallel and Distributed Processing Symposium* (IPDPS), 1091-1100, 2016.

[42] V. Varadarajan, T. Kooburat, B. Farley, T. Ristenpart, and M. Swift, "Resource-freeing attacks: improve your cloud performance (at your neighbor's expense)", *ACM Conference on Computer and Communications Security* (CCS), 2012.

[43] S. Govindan, D. Wang, A. Sivasubramaniam, and B. Urgaonkar, "Leveraging stored energy for handling power emergencies in aggressively provisioned datacenters", *International Conf. on Architectural Support for Programming Languages and Operating Systems* (ASPLOS), 2012.

[44] Y. Chen, A. Das, W. Qin, A. Sivasubramaniam, Q. Wang, and N. Gautam, "Managing server energy and operational costs in hosting centers", *International Conference on Measurement and Modeling of Computer Systems* (SIGMETRICS), 2005.

[45] C. Hsu, Q. Deng, J. Mars, and L. Tang, "SmoothOperator: Reducing Power Fragmentation and Improving Power Utilization in Large-scale Datacenters", *International Conf. on Architectural Support for Programming Languages and Operating Systems (*ASPLOS), 2018.

[46] D. Wang, S. Govindan, A. Sivasubramaniam, A. Kansal, J. Liu, and B.M. Khessib, "Underprovisioning backup power infrastructure for datacenters", International *Conf. on Architectural Support for Programming Languages and Operating Systems* (ASPLOS), 2014.

[47] Q. Wu, Q. Deng, L. Ganesh, C. Hsu, Y. Jin, S. Kumar, J. Meza, and Y. Song, "Dynamo: Facebook's Data Center-Wide Power Management System", *ACM/IEEE 43rd Annual International Symposium on Computer Architecture* (ISCA), 2016.

[48] 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.

[49] C. Li, R. Wang, D. Qian, and T. Li, "Managing Server Clusters on Renewable Energy Mix", ACM Transactions on Autonomous and Adaptive Systems (TAAS), Volume 11, Issue 1, 2016.

[50] C. Li, R. Wang, N. Goswami, X. Li, T. Li, and D. Qian, "Chameleon: Adapting throughput server to time-varying green power budget using online learning", *International Symposium on Low Power Electronics and Design* (ISLPED), 2013

[51] B. Aksanli, E. Pettis, and T. Simunic, "Architecting Efficient Peak Power Shaving Using Batteries in Data Centers", *IEEE International Symposium on Modelling, Analysis and Simulation of Computer and Telecommunication Systems* (MASCOTS), 2013.

[52] L. Liu, C. Li, H. Sun, Y. Hu, J. Xin, N. Zheng, and T. Li, "Leveraging Heterogeneous Power for Improving Datacenter Efficiency and Resiliency", *IEEE Computer Architecture Letters*, 2015

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

Xiaofeng Hou is currently a PhD candidate in the Department of Computer Science and Engineering at Shanghai Jiao Tong University. Her research includes data center resource management, reliable and dependable systems, and emerging cloud platforms. She has received a Best Paper Award from ICCD in 2018.

Jinghang Yang was a senior undergraduate student at Shanghai Jiao Tong University, China in 2018. After receiving his BS degree, he went to the University of Boston for M.S study. His research interests include cloud computing and big data systems.
ss

Chao Li received his Ph.D. degree from the University of Florida in 2014. He is currently a tenure-track associate professor in the Department of Computer Science and Engineering, Shanghai Jiao Tong University. His research mainly focuses on computer architecture and emerging computing systems.

Wenli Zheng received his PhD degree in electrical and computer engineering from The Ohio State University in 2016. He is currently a tenure-track assistant professor at Shanghai Jiao Tong University. His research interests include computer architecture, data center power management, and data-driven computing.

Xiaoyao Liang received his PhD degree from Harvard University. He is a professor in the Department of Computer Science and Engineering, Shanghai Jiao Tong University. His research interests include system architectures, energy efficient and resilient microprocessor, hardware/software co-design etc. He has ample industry experience working as a senior architect or IC designer at companies like NVIDIA, Intel and IBM.

Minyi Guo is a Zhiyuan Chair Professor in the Department of Computer Science and Engineering at Shanghai Jiao Tong University, China. He received his PhD degree in computer science from the University of Tsukuba, Japan. His research interests include parallel and distributed computing, compiler optimizations, computer architecture, cloud computing and big data. Prof. Guo has more than 250 publications in major journals and international conferences in these areas. He is a Fellow of IEEE