

Exploiting a Thermal Side Channel for Power Attacks in Multi-Tenant Data Centers

Mohammad A. Islam
University of California, Riverside

Shaolei Ren
University of California, Riverside

Adam Wierman
California Institute of Technology

ABSTRACT

The power capacity of multi-tenant data centers is typically oversubscribed in order to increase the utilization of expensive power infrastructure. This practice can create dangerous situations and compromise data center availability if the designed power capacity is exceeded. This paper demonstrates that current safeguards are vulnerable to well-timed power attacks launched by malicious tenants (i.e., attackers). Further, we demonstrate that there is a physical side channel — a thermal side channel due to hot air recirculation — that contains information about the benign tenants’ runtime power usage and can enable a malicious tenant to time power attacks effectively. In particular, we design a state-augmented Kalman filter to extract this information from the side channel and guide an attacker to use its maximum power at moments that coincide with the benign tenants’ high power demand, thus overloading the shared power capacity. Our experimental results show that an attacker can capture 54% of all attack opportunities, significantly compromising the data center availability. Finally, we discuss a set of possible defense strategies to safeguard the data center infrastructure against power attacks.

KEYWORDS

Data center; power attack; thermal side channel

1 INTRODUCTION

The explosion of cloud computing and the Internet of Things has generated a huge demand for multi-tenant data centers (also called “colocation”), resulting in a double-digit annual growth rate [1]. There are already nearly 2,000 multi-tenant data centers in the U.S. alone, accounting for five times energy of Google-type data centers combined altogether [2, 3]. Unlike a multi-tenant cloud platform that offers virtual machines (VMs), a multi-tenant data center is a shared facility where multiple tenants co-locate their own *physical* servers and the data center operator only manages the non-IT infrastructure (e.g., power and cooling). It serves almost all industry sectors, including top-brand IT companies (e.g., Apple houses 25% of its servers in multi-tenant data centers [4]).

The growing demand for multi-tenant data centers has created an increasingly high pressure on their power infrastructure (e.g., uninterrupted power supply, or UPS), which is very costly to scale up due to the high availability requirement (e.g., 99.9+%) and already

approaches the capacity limit in many cases [5]. The capital expense for data center infrastructure is around U.S.\$10-25 for each watt delivered to the IT equipment, exceeding 1.5 times of the total energy cost over its lifespan [6–8].

As a result, maximizing the utilization of the existing infrastructure in order to defer and/or reduce the need for expansion is a key goal for data center operators. To accomplish this, operators of multi-tenant data centers typically oversubscribe their power infrastructure by selling power capacity to more tenants than can be supported, counting on tenants not to have peaks in their power consumption simultaneously [9]. The industry standard is to oversubscribe the power capacity by 120% (yielding 20% more revenue for the operator at no extra cost) [10, 11]. This is also a common practice in owner-operated data centers (e.g., Facebook [8]) for improving power capacity utilization, and recent research has begun to suggest even more aggressive oversubscription [12, 13].

Power oversubscription is a powerful tool for increasing utilization and reducing capital cost, but it can potentially create dangerous infrastructure vulnerabilities. In particular, the designed power capacity can be overloaded (a.k.a. power emergency) when the power demand of multiple tenants peaks simultaneously. While data center infrastructure can tolerate short-term spikes, prolonged overloads over several minutes will make circuit breakers trip and result in power outages that are costly and may take hours or days to recover from [14–17]. For example, Delta Airlines incurs a US\$150 million loss due to a 5-hour power outage in its data center [18].

Although infrastructure redundancy is common in data centers and can absorb some overloads, they are not as reliable as desired. In fact, power equipment failures have now topped cyber attacks and become the most common reason for data center outages [16]. More importantly, *such redundancy protection is lost during power emergencies*, which is extremely dangerous and increases the outage risk by 280+ times compared to a fully-redundant case [19]. In fact, according to the data center tier classification (a higher tier means a better availability and hence higher construction cost) [19, 20], even though power emergencies only occur and compromise redundancy protection for 5% of the time, the expected downtime for a Tier-IV data center can increase by nearly 14 times to a similar level as a Tier-II data center, effectively resulting in a capital loss of 50% for the data center operator (Sec. 2.3).

Given the danger of power emergencies, an owner-operated data center operator can apply various power capping techniques (e.g., throttling CPU as done by Facebook [8]) to eliminate power emergencies. However, *a multi-tenant data center operator cannot follow similar approaches since it does not have the ability to control tenants’ servers.* In particular, a power emergency may occur while all tenants are operating within their own subscribed power capacities due to the operator’s power oversubscription. In such cases, the data center operator cannot forcibly cut power supplies to tenants’ servers without violating the contract; thus multi-tenant

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CCS’17, Oct. 30–Nov. 3, 2017, Dallas, TX, USA.

© 2017 ACM. 978-1-4503-4946-8/17/10...\$15.00

DOI: 10.1145/3133956.3133994

data centers are more vulnerable to power emergencies than owner-operated data centers [17].

As a consequence, multi-tenant data center operators have taken alternative precautions. They typically impose contractual terms to restrict tenants’ “normal” power usage to be below a certain fraction of their subscribed capacities (e.g., 80%), only allowing tenants to make limited use of the full subscribed capacities. Non-compliant tenants may face involuntary power cuts and/or eviction [21, 22]. This effectively avoids most, if not all, severe power emergencies, enabling the operator to safely oversubscribe its power capacity with a reasonably low risk of (usually mild) emergencies [11, 23]. As such, despite the common power oversubscription, power supply to tenants’ servers has long been considered as *safe* in a multi-tenant data center [24].

Contributions of this paper. *This paper focuses on an emerging threat to data center availability — maliciously timed high power loads (i.e., power attacks) — and highlights that multi-tenant data centers are vulnerable to power attacks that create power emergencies if power infrastructure oversubscription is exploited.* In particular, we demonstrate that, through observation of a thermal side channel, a malicious tenant (i.e., attacker) can launch well-timed *power attacks* with a high chance of successfully creating power emergencies that can potentially bring down the data center facility.

More specifically, although power emergencies are almost non-existent under typical operation due to statistical multiplexing of the servers’ power usage across benign tenants, a malicious tenant (which can be a competitor of the target multi-tenant data center) can invalidate the anticipated multiplexing effects by intentionally increasing its own power load up to its subscribed capacity at moments that coincide with high aggregate power demand of the benign tenants. This can greatly increase the chance of overloading the shared power capacity, thus threatening the data center uptime and damaging the operator’s business image.

In order to create severe power emergencies, the attacker must precisely time its power attacks. This may seem impossible because the attacker cannot use its full subscribed capacity continuously or too frequently, which would lead the attacker to be easily discovered and evicted due to contractual violations. Further, the attacker does not have access to the operator’s power meters and does not know the aggregate power usage of benign tenants at runtime.

The key idea we exploit is that the physical co-location of tenants’ servers in a shared facility means the existence of an important side channel — a thermal side channel due to heat recirculation.

Concretely, almost all server power is converted into heat, and some of the hot air exiting the servers may recirculate and travel a few meters to other server racks, (due to the lack of heat containment [24] in many data centers as shown in Section 3.3.1), which impacts the inlet temperature of those other racks [25, 26]. Heat recirculation constitutes an important side channel that the attacker can exploit to estimate the power of nearby tenants sharing the same power infrastructure. Nonetheless, since servers housed in different racks have different impacts on the attacker’s server inlet temperature, detection of a high temperature does not necessarily mean a high aggregate power usage of benign tenants.

To exploit the thermal side channel for timing power attacks, we propose a novel model-based approach: the attacker can build an

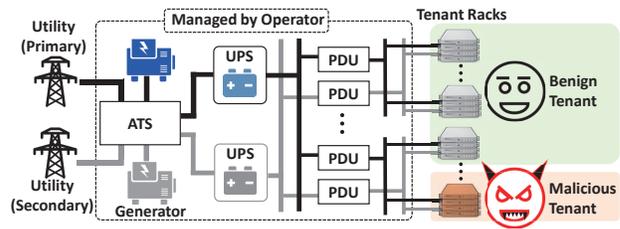


Figure 1: Tier-IV data center power infrastructure with 2N redundancy and dual-corded IT equipment.

estimated model for heat recirculation and then leverage a *state-augmented* Kalman filter to extract the hidden information about benign tenants’ power usage from the observed temperatures at its server inlets. By doing so, the attacker can control the timing of its power attacks without blindly or continuously using its maximum power: attacks are only launched when the aggregate power of benign tenants is also high. *Our trace-based experiments demonstrate that, with the aid of our proposed Kalman filter, an attacker can successfully capture 54% of all the attack opportunities with a precision rate of 53%, which significantly outperforms random attacks and represents state-of-the-art timing accuracy.* We also discuss possible defense strategies to safeguard the data center infrastructure, e.g., randomizing cooling system operation and early detection of malicious tenants (Sec. 5).

In conclusion, the key novelty of this paper is that it is the first study on power attacks in multi-tenant data centers by exploiting a thermal side channel. Our work is different from the existing data center security research that has mostly focused on *cyber* space, such as exhausting the IT resources (e.g., bandwidth via distributed denial of service, or DDoS, attacks [27, 28]) and co-residency attacks in the cyber domain (e.g., VM co-residency attacks [29, 30]). Moreover, in sharp contrast with the small but quickly expanding set of papers [12, 31, 32] that attempt to create power emergencies in an owner-operated data center, our work focuses on a multi-tenant setting and exploits a unique co-residency physical side channel — the thermal side channel due to heat recirculation — to launch *well-timed* power attacks.

2 IDENTIFYING POWER INFRASTRUCTURE VULNERABILITIES

This section highlights why and how power oversubscription happens in multi-tenant data centers. Additionally, it shows that if exploited by a malicious tenant through well-timed power attacks, power oversubscription can lead to emergencies, significantly compromising the data center availability.

2.1 Multi-tenant Power Infrastructure

A multi-tenant data center typically delivers protected power to tenants’ servers through multiple stages following a hierarchical topology. First, a UPS system takes utility power as its input and then outputs conditioned power to one or more power distribution units (PDUs). Next, each PDU steps down its input voltage and delivers power to a few tens of server racks at a suitable voltage.

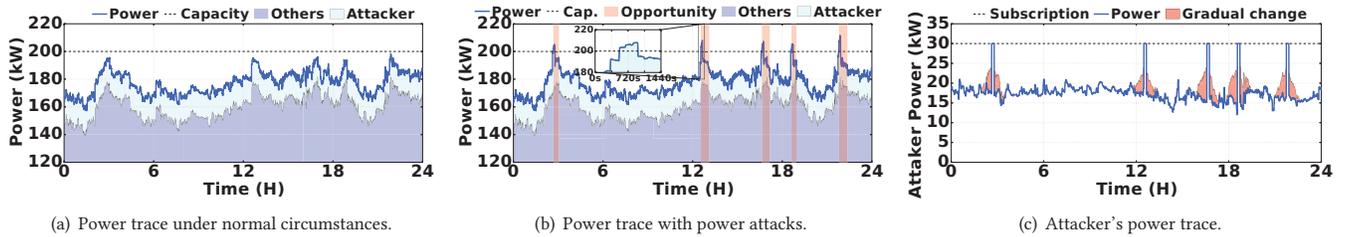


Figure 2: Infrastructure vulnerability to attacks. (a) Power emergencies are almost nonexistent when all tenants are benign. (b) Power emergencies can occur with power attacks. (c) The attacker meets its subscribed capacity constraint. The shaded part illustrates how the attacker can remain stealthy by reshaping its power demand when anticipating an attack opportunity.

Finally, each rack has a power strip (also called rack PDU) that supports a whole rack of servers. All the power equipment have circuit breakers, which will trip to prevent more serious consequences in case of a prolonged overload.

The power delivered to the IT equipment is also called *critical power*. Additionally, cooling system is needed to remove server heat, and its capacity is sized based on the critical power (i.e., cooling load). Thus, although data center capacity includes both power and cooling infrastructure capacities, it is often measured in the amount of total designed power capacity — total critical power supported by the power infrastructure subject to a certain availability requirement. In this paper, we follow this convention and use “(designed) power capacity” to refer to data center capacity. That is, *overloading the designed power capacity also implies overloading and stressing the designed cooling capacity*. Note that cooling system is connected to the utility substation through a separate path different from the servers.

To ensure a high infrastructure availability, redundancy is common in multi-tenant data centers. For example, Fig. 1 illustrates a fully-redundant Tier-IV facility, where the actually provisioned infrastructure capacity is twice as much as the designed power capacity to ensure an availability of 99.995+% [19, 33].

Data center capacity is leased to tenants on a per-rack basis according to the *designed* power capacity. Each tenant has multiple racks and needs to satisfy a per-rack power capacity constraint, while the operator is responsible for managing UPS/PDU units as well as the cooling system. While traditionally each centralized UPS unit has a capacity in the order of megawatt, many multi-tenant data centers have adopted modular construction by installing smaller UPS units (e.g., 100-200kW), each supporting one or a small number of PDUs. Thus, in a megawatt Tier-IV multi-tenant data center, there can exist several sets of 2N redundant infrastructures, each with a smaller designed capacity. Likewise, data center capacity is leased in a modular manner: only when the existing designed power capacity is fully leased will new capacity be built.

2.2 Vulnerability to Power Attacks

Due to its high capital expense (CapEx), power capacity is commonly oversubscribed by the data center operator, with an industry average oversubscription ratio of 120% [8, 23]. That is, the total power capacity leased to tenants is 120% of the capacity that is actually designed by the operator.

Oversubscribing the designed power capacity might result in *power emergencies*: the designed power capacity is overloaded when all the supported racks have their peak power usage simultaneously. Thus, the operator monitors each tenant’s power and typically imposes contractual terms to limit its normal usage to a fraction of the subscribed power capacity (e.g., 80%), while only allowing occasional and temporary usage of the full capacity [21, 22]. The contractual constraint can effectively make the tenants’ aggregate power demand stay well below the designed power capacity, thus achieving a high availability.

To illustrate this point, we show a 24-hour trace of power usage by four tenants in Fig. 2(a). The total designed power capacity is 200kW, but sold as 240kW because of the 120% oversubscription.¹ When all four tenants are benign, we see in Fig. 2(a) that power emergencies are almost nonexistent: there is no overload for the designed power capacity in our 24-hour snapshot. Indeed, even when a power emergency occurs due to coincident peak power usage of benign tenants, the overload is typically transient (because of the operator’s contractual constraint) and can be well absorbed by the power infrastructure itself [34].

In contrast, if a tenant is malicious, well-timed power attacks can successfully create prolonged power emergencies (e.g., overloading the designed capacity for several minutes). To see this point, we consider the same power trace as in Fig. 2(a), but inject power attacks by increasing the power usage of one tenant (i.e., attacker, which subscribes a total of 30kW power capacity) to its full capacity for 10 minutes whenever the designed capacity can be overloaded. The aggregate power demand and attacker’s power usage are shown in Fig. 2(b) and Fig. 2(c), respectively. In contrast to the benign case in Fig. 2(a), we see that five overloads of the designed power capacity occur over a 24-hour period in the presence of an attacker, while the attacker only uses its full power occasionally without continuously peaking its power or violating the operator’s contract [21]. In fact, even a benign tenant may have such usage patterns, but unlike malicious attacks, such benign peak power usage is not intentionally timed to create power emergencies and hence is much less harmful than malicious attacks (see Fig. 13 for a comparison between malicious attacks and random peaks).

The previous examples illustrate that the way that today’s multi-tenant data centers are managed is highly vulnerable: a malicious

¹More details of the power trace are provided in Section 4.1.

tenant can intentionally time its high power usage when the demand of benign tenants is also high, thus overloading the designed power capacity (shared by multiple tenants) much more often than otherwise would be.

2.3 Impact of Power Attacks

Data centers are classified into four tiers based on the degree of infrastructure redundancy in accordance with TIA-942 standard and the Uptime Institute certification [20, 33]. Next, we highlight that power emergencies created by malicious power attacks are very dangerous and significantly compromise the data center availability.

Tier-I. A basic Tier-I data center has no infrastructure redundancy: the actual provisioned capacity is the same as the designed capacity. Thus, it is cheaper to build (\$10/Watt capacity), but only has an availability of 99.671% which translates into an expected outage time of 28.80 hours per year [19, 35]. While power infrastructure can tolerate short-term spikes, prolonged overloads over a few minutes will alert the system and make the circuit breakers trip in order to prevent more catastrophic consequences (e.g., fire) [34]. See Fig. 18 in the appendix for the tripping time for a standard circuit breaker. Therefore, an overload of the designed capacity created by a successful power attack can easily bring down a Tier-I data center.

Tier-II/-III. A Tier-II/-III data center has “N+1” redundancy: if N primary non-IT units are needed for the designed capacity, then 1 additional redundant unit is also provisioned [19, 35]. Thus, overloading the designed capacity may not cause a data center outage, but will compromise the desired redundancy protection. For example, when any of the $N + 1$ units fails, overloading the designed capacity will bring down the remaining N infrastructure.

Tier-IV. A Tier-IV data center is fully $2N$ redundant: duplicating each needed non-IT unit, as illustrated in Fig. 1 [19, 35]. The redundant infrastructure may equally share the IT power loads with the primary infrastructure (“sharing” mode), or stand by and take over the loads when the primary infrastructure is overloaded or fails (“standby” mode) [36]. In either case, during an emergency that overloads the designed power capacity, such redundancy protection is lost: if with an emergency, a power outage can occur when either the primary or secondary infrastructure fails, but otherwise, it only occurs when *both* the primary and secondary infrastructures fail.

We now summarize the impact of power attacks in Table 1, by assuming that malicious power attacks result in emergencies (each lasting for 10 minutes) for 5% of the time. We first show the data center availability and corresponding expected outage time per year for each tier [19]. *The outage time only includes unplanned infrastructure failures*, while other types of outages, e.g., caused by human errors and cyber/network attacks, are excluded. While best operational practices may further improve availability, the availability value in Table 1 is representative for each tier based on real-world site measurement [19, 35].

Naturally, with power attacks, the expected outage time increases due to overloads of the designed capacity. For a Tier-1 data center, an overload of a few minutes will cause an outage as the circuit breakers will trip to prevent more serious consequences [34]. For a Tier-II/-III data center, we calculate the expected outage probability as “95% · (1 - p_a) + 5% · p_f ”, where p_a is the availability without

overloads and p_f is the failure rate of the redundant system. As redundancy increases the availability from $p_{a,I}$ to p_a (when there is no overload), we estimate $p_f = \frac{1-p_a}{1-p_{a,I}}$, where $p_{a,I}$ is primary system availability (using Tier-I availability value). For a fully-redundant Tier-IV data center, we assume that the primary and redundant systems are completely independent [19], each having a failure rate of $\sqrt{1-p_a}$ without overloads. Thus, with emergencies occurring for 5% of the time, the outage probability can be estimated as “95% · (1 - p_a) + 5% · $[2\sqrt{1-p_a} - (1-p_a)]$ ”. Then, we show the expected outage time with attacks per year, as well as the new availability values. A higher-tier data center is more costly to build, e.g., the capital expense for each watt of critical power for a Tier-IV data center is twice as much as a Tier-II data center [35]. Nonetheless, due to the increased outage time exceeding the tier standard, the intended tier classification may not apply anymore. Such tier downgrading essentially means a capital loss for the operator (i.e., higher cost for a lower tier), which is also shown in Table 1 based on the power capacity cost data in [35]. It will also damage the operator’s business image in the long term and result in a customer turnover.

In addition, power attacks also lead to increased outage costs borne by affected tenants (compared to the no-attack case). For example, even a power outage in a single data center can cost millions of dollars, as exemplified by the recent British Airways data center outage [37]. Although application-level redundancy across geo-distributed data centers may retain service continuity during outages in a single location, the workload performance of affected tenants can be significantly degraded due to traffic re-routing and migration [38, 39]. We estimate the average outage cost per sqft per minute based on [16], excluding service losses due to recovery after an outage. The outage costs are 0.033, 0.1073, 0.7783 and 0.93 (all in “\$/sqft/min”), for Tier-I to Tier-IV data centers, respectively. The total outage cost increase is shown in Table 1, which is even higher than the operator’s amortized capital loss.

In conclusion, even if emergencies only occur for 5% of the time due to power attacks, data center availability is significantly compromised, resulting in a huge financial loss for both the operator and benign tenants.

3 EXPLOITING A THERMAL SIDE CHANNEL

The previous section highlighted the danger of maliciously timed power attacks that can compromise long-term data center availability. In this section, we exploit a thermal side channel to estimate the aggregate power usage of benign tenants and, thus, guide an attacker to time its attacks against the shared power infrastructure.

3.1 Threat Model

We consider an oversubscribed multi-tenant data center where a malicious tenant, i.e., attacker, houses physical servers and shares a designed power capacity of C with several other benign tenants. The attacker’s servers can be divided into groups and deployed under multiple accounts in different locations inside the target data center (to better estimate the power consumption of nearby benign tenants as shown in Sec. 3.4.1). While it may be possible that the attacker hides advanced weapons/bombs in its modified servers to physically damage the facility, such attacks are orthogonal to our

Table 1: Estimated impact of power emergencies (5% of the time) on a 1MW-10,000sqft data center.

Classification	Specification	Outage (hours/Yr)	Outage w/ Attack (hours/Yr)	Increased Outage Cost (mill. \$/Yr)	Capital Loss (mill. \$/Yr)	Total Cost (mill. \$/Yr)
Tier-I (Availability: 99.671%)	No redundancy	28.82	465.36 (Availability: 94.688%)	8.57	NA	8.57
Tier-II (Availability: 99.741%)	N+1 redundancy (generator/UPS/chiller)	22.69	366.36 (Availability: 95.818%)	22.11	0.1 (9+%↓)	22.22
Tier-III (Availability: 99.982%)	N+1 redundancy (all non-IT equipment)	1.58	25.46 (Availability: 99.709%)	11.15	1.0 (50%↓)	12.15
Tier-IV (Availability: 99.995%)	2N redundancy (all non-IT equipment)	0.44	6.59 (Availability: 99.925%)	3.42	1.1 (50%↓)	4.52

work. Instead, we focus on an unexplored threat model: an attacker aims to compromise the data center infrastructure availability by maliciously timing its peak power usage. That is, *the attacker behaves normally as other benign tenants, except for that it launches power attacks to create power emergencies by intentionally using its full subscribed power capacity when it anticipates a high aggregate power of benign tenants.*

We consider an attack successful if “ $p_a + p_b \geq C$ ” is satisfied over a continuous time window of at least L minutes ($L = 5$ in our default case and is enough to trip an overloaded circuit breaker [34]), where p_a is the attacker’s power and p_b is the aggregate power of benign tenants. Accordingly, we say that there is an *attack opportunity* if a successful attack can be possibly launched by the attacker, regardless of whether an attack actually occurs.

- **What the attacker can do.** We assume that the attacker knows the shared power capacity (as advertised by the operator) and can subscribe to a certain amount of capacity at a fairly low price (e.g., monthly rate of U.S.\$150/kW [40]). Then, when an attack opportunity arises, the attacker can generate malicious power loads almost instantly by running simple CPU-intensive workloads.² As servers are merely used to launch power attacks, the attacker does not need to run any useful workloads and can install any low-cost (even second-hand) high-power servers in its racks. In order to stay stealthy, the attacker can gradually increase power and also reshape its power demand when it anticipates a power attack opportunity, as illustrated in solid color in Fig. 2(c). Further, we assume that the attacker conceals temperature sensors at its server inlets and can perform computational fluid dynamics (CFD) analysis, which is a standard tool for modeling data center heat recirculation and easy to use for anyone with a good knowledge of data center operation (see Autodesk tutorial [41]).

- **What the attacker cannot do.** In our model, the attacker cannot hide destructive weapons inside its servers for attacks, which may not even pass the move-in inspection by the operator and can be held legally liable. Neither can the attacker modify its off-the-shelf servers to, for example, generate transient power spikes/pulses. These spikes/pulses may trip the attacker’s own rack circuit breaker and/or be detected by the operator’s power monitoring system.

²If some benign tenants offer web-based services open to the public, the attacker may also remotely send more requests to these benign tenants’ servers to increase their power consumption when it detects an attack opportunity.

Given the attacker’s access to the target data center, there may exist other attack opportunities to bring down a data center, such as congesting the shared bandwidth, which are complementary to our focus on attacking the shared non-IT infrastructure and compromising its designed availability. Moreover, we do not consider remotely hacking the data center infrastructures or manually tampering with the power infrastructures (all tenants’ visits to a multi-tenant data center are closely monitored and escorted). These may be possible, but are orthogonal to our study.

- **Who can be the attacker?** The attacker’s cost (i.e., server cost plus data center leasing cost) is only a small fraction of the benign tenants’ financial loss or operator’s capital loss (between 1.44% and 15.88%, Sec. 4.2.2), thus providing a sufficient motivation for the attacker. For example, the attacker can be a competitor of the target multi-tenant data center, which not only results in the victim’s capital loss but also significantly damages its business image. Note that power outages result in power cut for both benign tenants and the attacker (which does not run useful workloads), and these are what the attacker is aiming to create.

To summarize, we focus on malicious power attacks to overload the shared power infrastructure in a multi-tenant data center and compromise its availability. Towards this end, an attacker intentionally creates power emergencies by using its peak power when the benign tenants’ aggregate power demand is high. Meanwhile, the attacker’s power consumption still meets the operator’s contractual constraint.

3.2 The Need for a Side Channel

As illustrated in Fig. 2(b), attack opportunities exist intermittently due to the fluctuation of benign tenants’ power usage. Thus, a key question is: *how does the attacker detect an attack opportunity?*

Naturally, an attack opportunity arises when the aggregate power of benign tenants is sufficiently high. But, the benign tenants’ power usage is only known to themselves and to the data center operator (through power meters) – not to the malicious tenant.

A naive attacker may try to always use the maximum power allowed by its subscribed capacity in order to capture all attack opportunities. But, this is not realistic since the attacker may face power supply cut (due to violation of contractual terms) and be evicted [21, 22]. Similarly, blindly or randomly launching attacks at random times is not likely to be successful (Fig. 13 in Sec. 4.2.2).

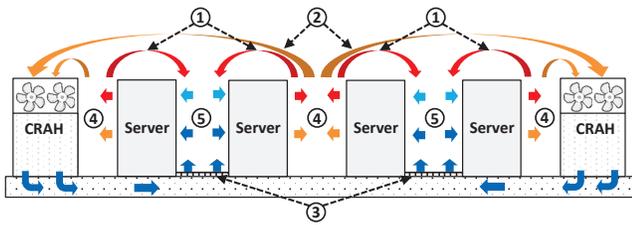


Figure 3: Cooling system overview. (1) Hot recirculated air. (2) Return air. (3) Perforated tile. (4) Hot aisle. (5) Cold aisle.

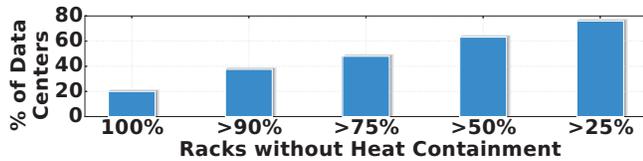


Figure 4: Adoption of heat containment techniques [24].

Another naive strategy for the attacker would be to simply select a *coarse* opportunity window to launch attacks. For example, the attacker may choose *peak* hours. Nonetheless, the multiplexing of independent tenants that run diverse workloads means that the aggregate peak power usage can occur more randomly and *outside* of expected peak hours. Alternatively, with dual power supply in a Tier-IV data center illustrated in Fig. 1, the attacker can detect the loss of infrastructure redundancy protection when seeing that only one cord is supplying power (which may take several hours to correct); then, it can launch power attacks in order to bring down the data center. But, such dual power supply may not be available in all data centers (especially lower-tier data centers [19]).

Even though a coarse opportunity window exists (e.g., peak hours occur regularly or failure of the secondary infrastructure is detected) and helps the attacker locate the attack opportunities within a smaller time frame, the actual attack opportunity is intermittent and may not last throughout the entire coarse window, as shown in Fig. 2(b). Thus, the attacker needs a *precise* timing in order to launch successful attacks with a higher chance. For this reason, side channels that leak (even noisy) information about the benign tenants' power usage at runtime are crucial for the attacker.

3.3 A Thermal Side Channel

An important observation is that the co-residency of the attacker and benign tenants in a shared physical space means that a thermal side channel exists. To see why, let us look at how the cooling system works in a typical multi-tenant data center.

3.3.1 Cooling System Overview. A cooling system is essential for conditioning the server inlet temperature (between 65°F and 81°F [42]) and maintaining data center uptime [43]. Most multi-tenant data centers, especially medium and large ones, adopt a *raised-floor* design and use computer room air handlers (CRAHs) in conjunction with outdoor chillers to deliver cold air to server racks [44, 45]. Smaller data centers often rely on computer room air conditioning (CRAC) units, which use compressors to produce

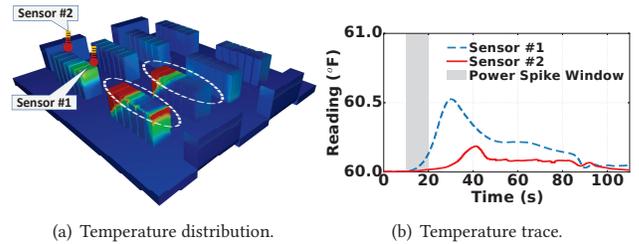


Figure 5: CFD simulation result. (a) Temperature distribution after 10 seconds of a 10-second 60kW power spike at the circled racks. (b) Temperature trace at select sensors.

cold air. For both types of systems, the indoor part is similar and illustrated in Fig. 3.

Cold air is first delivered by the CRAHs to the underfloor plenum at a regulated pressure greater than the room air pressure [46]. The air pressure difference pushes the cold air upwards through perforated tiles. After entering the servers through *server inlets*, the cold air absorbs the server heat and then exits the servers.

The CRAH controls the volume of its air supply to maintain a target air pressure at select sensor locations underneath the floor. Further, the opening area of perforated tiles is often manually set and changed only when server rack layout/power density is changed [43, 47]. As such, there is not much frequent variation in the flow rate of cold air entering the data center room.

For delivering cold air dynamically to accommodate variable demands and improving efficiency, heat containment (e.g., seal cold/hot aisles to decrease heat recirculation) is needed [48]. Nonetheless, heat containment needs a high level of homogeneity in server rack layout, and some tenants may be concerned with the potential risks (e.g., fire safety) [49, 50]. Thus, as illustrated in Fig. 3, many multi-tenant data centers rely on an open airflow path to serve multiple tenants. This is also confirmed by a recent Uptime Institute survey covering 1,000+ large/medium data centers [24] which, as plotted in Fig. 4, shows that nearly 80% data centers have at least 25% of racks without heat containment and that 20% data centers do not have any heat containment at all.

In our experiment (Fig. 14(c)), we will investigate how different levels of heat containment will affect the timing accuracy of power attacks.

3.3.2 Heat Recirculation. Although most hot air directly returns to the CRAHs to be cooled down, some hot air can travel a few meters to other server racks in the shared open space and impact their inlet air temperature [25, 26, 51]. To better illustrate this phenomenon (called *heat recirculation*), we run industry-grade CFD simulations to model the data center airflow [52]. With all the servers at deep sleep states consuming nearly zero power, we generate a 10-second 60kW power load evenly distributed among 12 server racks (marked with circles) in Fig. 5(a). Ten seconds after the power spike, the data center temperature distribution is shown in Fig. 5(a), where blue and red surfaces represent low and high

temperatures, respectively. It can be clearly seen that the temperature of nearby racks is affected by the power spike. The detailed temperature changes at two select sensor locations are also shown in Fig. 5(b). We also show the breakdown of temperature readings monitored at sensor #1 in Fig. 20 (Appendix C).

Heat recirculation is generally undesirable for efficiency reasons [25, 51]; our work shows that it is undesirable for security reasons too, since it constitutes a thermal side channel that an attacker can use to launch well-timed power attacks. Concretely, the attacker’s server inlet temperature contains some, albeit not accurate, information of benign tenants’ power usage: if a server of a benign tenant consumes more power, it will result in a higher temperature increase at the attacker’s server inlets.

3.4 Estimating Benign Tenants’ Power from a Thermal Side Channel

Given the impact of the benign tenants’ power usage at the attacker’s server inlet temperature, the attacker may use this information to obtain (noisy) estimates of the aggregate power usage and launch well-timed power attacks.

An intuitive, but naive, approach is to launch power attacks based on a temperature threshold (which we call temperature-based power attack): attack when the temperature reading is higher than a threshold. Nonetheless, temperature-based power attacks are hardly better than launching attacks at random times.

To illustrate this point, we run CFD analysis (details in Section 4.1) and present a snapshot in Fig. 21 in the appendix. In our experiment, the attacker launches a 10-minute attack whenever its average temperature reading exceeds 76.3°F for at least 1 minute, but the snapshot shows that all attacks are unsuccessful. We further vary the temperature threshold for power attacks and show the result in Fig. 6(a). As expected, with a lower temperature threshold, the attacker attacks more frequently (e.g., 45+% of the times given a temperature threshold of 74°F) and can capture more attack opportunities, but the precision (i.e., the percentage of successful attacks among all the launched attacks) still remains very low. In practice, the attacker cannot use its full capacity too frequently due to contractual constraints. Thus, for practical cases of interest (e.g., launching attacks for no more than 10% of the times), the attacker can hardly capture any attack opportunities. We also consider power attacks based on the maximum temperature reading, and similar results are shown in Fig. 6(b).

The reason temperature-based power attacks have a poor detection of attack opportunities is that heat recirculation is spatially *non-uniform* (i.e., more significant among racks closer to each other), and hence different servers can result in drastically different temperature impacts on the attacker’s temperature sensors. Moreover, the attacker’s own power usage (as well as noise) also greatly impacts the temperature readings. Thus, temperature does not accurately reflect the benign tenants’ aggregate power usage, motivating us to study alternative approaches to making a better use of the prominent thermal side channel.

3.4.1 Modeling Heat Recirculation. As heat recirculation is spatially *non-uniform*, the same server, if placed in different racks, can have very different impacts on the attacker’s server inlet temperature. Thus, to better estimate the benign tenants’ power usage,

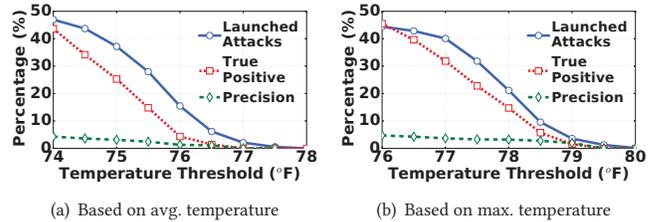


Figure 6: Summary of temperature-based power attacks. The line “Launched Attacks” represents the fraction of time power attacks are launched.

the attacker needs to further attribute its server inlet temperature increase to different servers. Such information can be extracted by the attacker with the help of a heat recirculation model, which relates a server’s power usage to the inlet temperature increase at the attacker’s servers. In what follows, as proposed in [25, 53], we present a simple yet accurate linear model of heat recirculation and quantify how an individual server’s power usage affects the attacker’s server inlet temperature.

Note that the actual heat recirculation model is unknown to the attacker; instead, *the attacker only has limited and imprecise knowledge of how heat recirculates in the data center*, which can deviate significantly from the actual process (Sec. 3.4.2 and Fig. 8). But, our experiments in Sec. 4.2 show that even imprecise knowledge of the heat recirculation model can assist the attacker time its power attacks with a high accuracy.

We consider a discrete time-slotted model, where the attacker has M sensors (indexed by $m = 1, 2, \dots, M$) and reads its temperature sensors once every time slot (e.g., every 10 seconds). There are N servers (indexed by $n = 1, 2, \dots, N$) owned by benign tenants. Meanwhile, the attacker owns J servers indexed by $n = N + 1, N + 2, \dots, N + J$. We denote the (average) power consumption of server n during time slot t as $p_n(t)$.

The attacker’s temperature sensor reading can be affected by a server’s power over the previous K time slots, because it takes time for hot air generated by a server to travel to the attacker’s server inlet (e.g., up to 100 seconds in our CFD simulations) [25]. Prior research [25, 53] has shown that, given a particular airflow pattern, the heat recirculation process can be modeled as a finite-response *linear* time-invariant system whose inputs and outputs are a server’s power usage and the temperature increase at a sensor, respectively.

Concretely, the cumulative temperature increase at sensor m caused by server n at time t can be expressed as $\Delta T_{m,n}(t) = p_n(t) * h_{m,n}(t) = \sum_{\tau=0}^{K-1} p_n(t-\tau) \cdot h_{m,n}(\tau)$, where “*” is the convolution operator and $h_{m,n}(t)$ is the system impulse response function (i.e., $h_{m,n}(t)$ denotes the temperature increase at sensor m at time t if server n consumes a unit power at time 0). Note that $h_{m,n}(t) = 0$ for $t < 0$ (due to system causality) and $t \geq K$ (since the hot air generated at a server only contributes to the sensor temperature increase for up to K time slots).

Next, we sum up the temperature impact caused by all the servers in the data center and express the m -th temperature sensor reading

at time t as

$$T_m(t) = T_{sup}(t) + \sum_{n=1}^{N+J} \sum_{\tau=0}^{K-1} p_n(t-\tau) \cdot h_{m,n}(\tau) + r_m(t), \quad (1)$$

where $T_{sup}(t)$ is the supply air temperature and $r_m(t)$ is the noise capturing random disturbances. For notational convenience we use $\vec{p}_{b,t} = \{p_1(t), \dots, p_N(t)\}$ to denote the vector of the power usage for benign tenants' N servers at time t . We also use a column vector $x_t = [\vec{p}_{b,t}, \vec{p}_{b,t-1}, \dots, \vec{p}_{b,t-K+1}]^T$, where "T" denotes the transpose, to include all the benign tenants' power usage values over the past K time slots. Similarly, for the attacker, we denote $\vec{p}_{a,t} = \{p_{N+1}(t), \dots, p_{N+J}(t)\}$ and use $y_t = [\vec{p}_{a,t}, \vec{p}_{a,t-1}, \dots, \vec{p}_{a,t-K+1}]^T$.

Next, we rewrite the model in (1) as follows

$$z_t = T_t - T_{sup}(t) \cdot \mathbf{I} - \mathbf{H}_a y_t = \mathbf{H}_b x_t + r_t, \quad (2)$$

where $T_t = [T_1(t), \dots, T_M(t)]^T$ is the vector of temperature readings, $\mathbf{I} = [1, 1, \dots, 1]^T$ is an $N \times 1$ identity vector, $r_t = [r_1(t), \dots, r_M(t)]^T$, and \mathbf{H}_a and \mathbf{H}_b are *heat recirculation matrices* containing all the system impulse functions that relate server power of the attacker and the benign tenants to the temperature increase at the attacker's sensors, respectively. In particular, the m -th row of \mathbf{H}_a is $[h_{m,N+1}(0), \dots, h_{m,N+J}(0), \dots, h_{m,N+1}(K-1), \dots, h_{m,N+J}(K-1)]$, while the m -th row of \mathbf{H}_b is $[h_{m,1}(0), \dots, h_{m,N}(0), \dots, h_{m,1}(K-1), \dots, h_{m,N}(K-1)]$.

3.4.2 A State-Augmented Kalman Filter. Kalman filters are a classic tool to estimate hidden states from noisy observations in many applications, such as power grid state estimation and aircraft control [54]. Here, we apply a Kalman filter to estimate benign tenants' runtime power usage, which is not directly observable but is contained in the thermal side channel.

Design of a Kalman filter. The observation model can be specified using the heat recirculation model in (2). As the current temperature reading is affected by the servers' power usage over the past K time slots, we use $x_t = [\vec{p}_{b,t}, \vec{p}_{b,t-1}, \dots, \vec{p}_{b,t-K+1}]^T$ as the augmented state. We also view $z_t = T_t - T_{sup}(t) \cdot \mathbf{I} - \mathbf{H}_a y_t$ as the equivalent observation (or measurement), because the $T_{sup}(t)$ is known (e.g., by placing an additional sensor at the perforated tile) and $\mathbf{H}_a y_t$ is the temperature increase due to the attacker's own power usage that is known to itself.

In addition, the attacker needs a *process* model to characterize the dynamics of benign tenants' power usage (i.e., state) over time, which is unknown to the attacker. Thus, for simplicity, the attacker assumes that the benign tenant's server power is driven by a noise process, i.e., $p_n(t+1) = p_n(t) + q_{n,t}$, where $q_{n,t}$ is the random noise. Thus, the process model can be written as

$$x_{t+1} = \mathbf{F}x_t + q_t. \quad (3)$$

In the model, $q_t = [q_{1,t}, q_{2,t}, \dots, q_{N,t}, 0, \dots, 0]$ is a $NK \times 1$ column vector with \mathbf{Q} being its covariance matrix, and $\mathbf{F} = [\mathbf{I}_{N \times N}, \mathbf{0}_{N \times N(K-1)}; \mathbf{I}_{N(K-1) \times N(K-1)}, \mathbf{0}_{N(K-1) \times N}]$ is a $NK \times NK$ matrix governing the state transition, where $\mathbf{I}_{n \times n}$ is an $n \times n$ diagonal matrix with 1 along the diagonal and 0 in all other entries and $\mathbf{0}_{m \times n}$ is an $m \times n$ zero matrix.

The thermal side channel is then fully characterized by combining the observation model in (2) and process model in (3). Thus, the

attacker can apply a Kalman filter to estimate x_t , which includes the benign tenants' power $\vec{p}_{b,t} = \{p_1(t), \dots, p_N(t)\}$ at time t .

Denoting $\hat{x}_{t|t-1}$ as an estimate of x at time t given observations up to time $t-1$ and \mathbf{R} as the covariance matrix of measurement noise, we show the key steps in a Kalman filter [54] as follows.

$$\begin{aligned} \text{Predict:} \quad \hat{x}_{t|t-1} &= \mathbf{F}\hat{x}_{t-1|t-1} \\ \mathbf{P}_{t|t-1} &= \mathbf{F}\mathbf{P}_{t-1|t-1}\mathbf{F}^T + \mathbf{Q} \\ \text{Update:} \quad u_t &= z_t - \mathbf{H}_b\hat{x}_{t|t-1} \\ \mathbf{S}_t &= \mathbf{H}_b\mathbf{P}_{t|t-1}\mathbf{H}_b^T + \mathbf{R} \\ \mathbf{G}_t &= \mathbf{P}_{t|t-1}\mathbf{H}_b^T\mathbf{S}_t^{-1} \\ \hat{x}_{t|t} &= \hat{x}_{t|t-1} + \mathbf{G}_t u_t \\ \mathbf{P}_{t|t} &= (\mathbf{I} - \mathbf{G}_t\mathbf{H}_b)\mathbf{P}_{t|t-1} \end{aligned}$$

Even if the supply temperature $T_{sup}(t)$ is unknown, we can append it after the power state and update the estimation procedure accordingly.

Practical considerations. Applying the Kalman filter above to estimate benign tenants' server-level power usage has two main challenges. First, it can be highly inaccurate as well as computationally expensive to estimate a large number of N hidden states, each representing the power usage of one server. Second, to estimate hundreds of hidden states based on the model in (1), the attacker needs to know a large heat recirculation matrix \mathbf{H}_b , i.e., $M \cdot N$ system impulse response functions $h_{m,n}(t)$.

To address these challenges, we propose to estimate benign tenants' power usage on a virtual *zonal* basis. Specifically, the attacker divides the target data center into multiple virtual zones (each containing one or more tenants) and estimates the power for each zone of racks as a single entity, rather than for each individual server. In fact, estimating zone-level power usage already suffices, because the attacker only needs to know the *aggregate* power usage of benign tenants.

To construct the zone-level heat recirculation matrix, the attacker can visit the data center (as any tenant can) and visually inspect its layout. Then, following the industry practice and as described in Section 4.1, the attacker can perform CFD analysis to construct a zone-level heat recirculation model with the assumption that all the servers in one zone yield the same temperature increase impact on the attacker's sensors.

Naturally, the zone-level heat recirculation model only *approximates* the detailed server-level model (1), and the attacker cannot exactly know the data center layout from a visual inspection. Thus, the attacker only has an estimate of the actual heat recirculation model. Despite this limitation, we show in Section 4.2 that the attacker can still estimate the benign tenants' aggregate power usage with high accuracy (e.g., only 3% error on average), capturing 54% of the attack opportunities.

3.5 Attack Strategy

In a typical multi-tenant data center, a tenant is allowed to use power up to $\alpha \cdot C_t$ continuously, where C_t is the power capacity subscribed by the tenant and $\alpha < 1$ is the threshold (usually 80%) set by the operator [21, 22]. A tenant can also use its full capacity C_t occasionally, but continuously using it can result in an involuntary power cut and eviction. We now discuss how the attacker can make

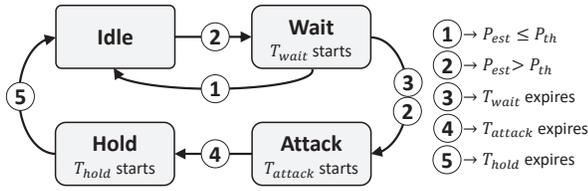


Figure 7: Finite state machine of our attack strategy. P_{est} is the attacker’s estimated aggregate power demand (including its own), and P_{th} is the attack triggering threshold.

use of the estimation procedure above in order to time its attacks while meeting the operator’s contractual constraint.

We consider a simple strategy where the attacker keeps on using its maximum power for a fixed time of T_{attack} , when it *anticipates* a high aggregate power usage of $P_{est} \geq P_{th}$ (called triggering threshold). The triggering threshold P_{th} is an important choice parameter for the attacker: the smaller P_{th} , the more attacks. Before launching an attack, the attacker should wait for its estimate of benign tenants’ usage to stay high for some time T_{wait} , in order to reduce unsuccessful attacks when the estimate of benign tenants’ power is only transiently high. Depending on the operator’s contractual constraint, we also set a hold time of T_{hold} before the attacker launches its next attack and impose a constraint on its triggering threshold $P_{th} \geq \hat{P}_{th}$. In our experiments (Sec. 4.2.2), we will vary how long and how frequently the attacker is allowed to fully use its subscribed capacity.

We illustrate our attack strategy using a finite state machine in Fig. 7. More advanced strategies are left as future work, as the current one is already quite effective.

4 EXPERIMENTAL EVALUATION

To demonstrate the danger of power attacks by malicious tenants, we evaluate how well the attacker can detect attack opportunities based on the thermal side channel. Our experimental results highlight that, with the aid of a Kalman filter and by launching attacks no more than 10% of the time, the attacker can successfully detect 54% of all attack opportunities (i.e., true positive rate) with a precision of 53%.

Although these values may vary depending on the specific settings, our results offer an important support to the broad implication: *the attacker can extract useful information about benign tenants’ runtime power usage from the thermal side channel and launch well-timed successful power attacks against the power infrastructure.*

4.1 Methodology

Because of the destructive nature of power attacks and the practical difficulty in accessing mission-critical data center facilities, we use an industry-grade simulator, Autodesk CFD [52], to perform CFD analysis and simulate heat recirculations driven by a real-world workload trace. The accuracy of CFD analysis has been well validated [25, 53], and many data centers, including Google [55], use CFD analysis to predict temperature distributions [25, 46, 53]. Thus, before any demonstration in industrial multi-tenant data centers is planned, the CFD-based simulation provides us with an

important understanding of the possibility and danger of power attacks timed through a thermal side channel. Our default settings are described below.

Data center layout. We consider a modular infrastructure design where a large data center is constructed using multiple independent sets of non-IT infrastructures, each having a smaller designed capacity. Specifically, the total designed capacity under consideration is 200kW and, according to the industry average [23], oversubscribed by 120%. We follow the design by HP Labs [46] and show the indoor part of our considered data center space in Fig. 19 (Appendix B). To get an idea of the heat recirculation process, the attacker divides the shared data center space into four different virtual zones (three for benign tenants and one for the attacker), while we note that the attacker’s zone division is not unique. Zones 1 and 2 have 12 server racks each. Zone 3 has 18 server racks, while the attacker occupies the 4th zone with 6 racks. Each rack has 20 servers and a power capacity of 5kW. There are four CRAH units that supply cold air to servers through perforated tiles.

CFD analysis. We port our data center layout shown in Fig. 19 into Autodesk CFD to quantify the heat recirculation process [52]. The physical components, such as servers, racks, raised floor and CRAH, are designed in Autodesk Inventor based on its data center simulation guideline [41]. Nonetheless, as CFD is computationally prohibitive, it cannot be used for simulations with month-long power traces. Thus, to calculate the attacker’s temperature, we follow the literature [25, 53] and use the server-level heat recirculation model in (1), where the system impulse response function $h_{mn}(t)$ is derived by generating a power spike over one time slot (10 seconds) for server n and getting the temperature at sensor m through the CFD analysis on Autodesk. This process is repeated for all the servers and sensors. The accuracy of the linear model in (1) has been extensively validated against real system implementations [25, 46, 51, 53]. Thus, the model has been widely applied to guide temperature-constrained runtime resource management [25, 51, 53]. Here, we use it for a new purpose — assisting the attacker with timing its power attacks.

Power trace. We collect a representative composition of four different power traces for the four virtual zones (Fig. 19), following the practice of prior studies [12, 17]. Specifically, two power traces are collected from Facebook and Baidu production clusters [8, 15] and used for virtual zones 1 and 2, respectively. We also collect two request-level batch workload traces (SHARCNET and RICC logs) from [56, 57], and, based on the power model validated against real systems [6], convert them into the power usage of the third virtual zone and the attacker. All the power usage are scaled to have an average utilization of 75% (for the 3 virtual zones) and 60% (for the attacker), normalized with respect to the tenant’s subscribed capacity. Fig. 9 shows a 24-hour snapshot of the synthetic aggregate power trace, which has a consistent pattern with real measurements [8, 15]. We also evaluate the Kalman filter performance and attack success rate on an alternative set of power traces in Appendix F.

Attacker. The attacker has 6 racks in one row as illustrated in Fig. 19. It has six sensors placed evenly along the top of its six racks, and reads the sensors once every 10 seconds (one time slot). The attacker’s sensor noise includes two parts: random disturbance/random noise following a Gaussian distribution $\mathcal{N}(0, 0.5)$

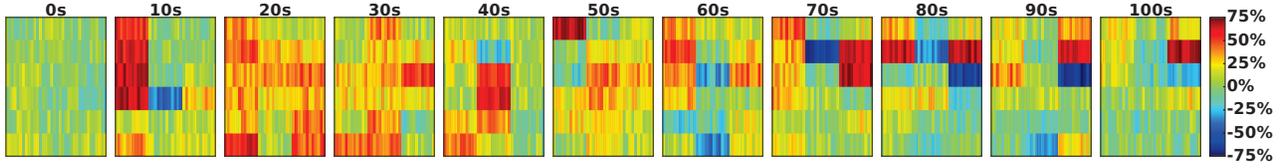


Figure 8: Error in the attacker’s knowledge of heat recirculation matrix H_b , normalized to the true value.

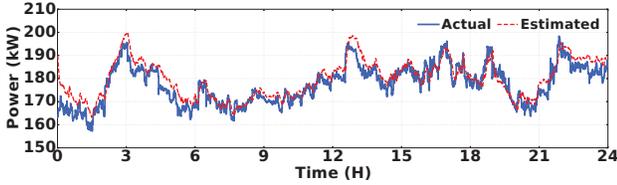


Figure 9: A snapshot of the actual and estimated power.

with a unit of $^{\circ}\text{F}$, and additional noise modeled as a variable that has a mean of 0.5°F and scales proportionally with the power trace in [15] (capturing the impact of servers that are served by other infrastructures but housed in the same room). Following the strategy in Sec. 3.5, after detecting an attack opportunity and waiting for $T_{wait} = 1$ minute, the attacker increases its power to the full capacity for $T_{attack} = 10$ minutes. By default, the attacker does not attack consecutively or more than 10% of the time each day, and sets $T_{hold} = 10$ minutes.

Note that if available, a coarse timing (e.g., daily peak hours, see Sec. 3.2) may help the attacker focus on a narrower time frame for attacks, but it is still inadequate due to the short duration of intermittent attack opportunities. In contrast, we focus on fine-grained *precise* timing by exploiting a thermal side channel, on top of the complementary coarse timing.

4.2 Evaluation Results

Our evaluation results highlight that our proposed Kalman filter can extract reasonably accurate information about benign tenants’ power usage and guide the attacker to launch successful attacks.

4.2.1 Kalman Filter Performance. The attacker constructs a zone-level heat recirculation matrix for its Kalman filter (Section 3.4.2) and hence, only has an inaccurate knowledge of the actual heat recirculation matrix H_b in (2). Given this limitation, let us first examine the Kalman filter performance.

In our experiment, we consider three zones for the benign tenants as illustrated in Fig. 19. We show the attacker’s estimate of zone-based temperature increase impact at one of its sensors in Fig. 22 (Appendix C). Further, we show in Fig. 8 the attacker’s error normalized with respect to the true heat recirculation matrix H_b , i.e., the values of $\frac{\hat{h}_{m,n}(t) - h_{m,n}(t)}{h_{m,n}(t)}$, where $\hat{h}_{m,n}(t)$ is the value generated by the attacker’s zone-based model. Each heat map indicates the normalized errors for one time slot. It has six rows and 840 columns, corresponding to the attacker’s 6 sensors and benign

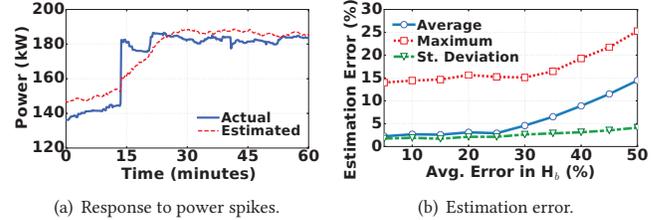


Figure 10: Robustness of Kalman filter performance. (a) The Kalman filter response to large power spikes. (b) Power estimation error versus error in the attacker’s knowledge of heat recirculation matrix.

tenants’ 840 servers, respectively. The average normalized error is 20%, while the maximum error is $\pm 75\%$.

Next, we show a 24-hour snapshot of the actual and estimated aggregate power in Fig. 9. While estimation errors can be large at certain times, the attacker’s estimate generally follows the same pattern of the actual power.

We now examine the Kalman filter robustness. The process model (3) assumes that the benign tenants’ power is driven by a noise, but this may not hold in practice. Thus, we create an artificial large power spike (unlikely in practice) and see how the filter responds. It can be observed from Fig. 10(a) that the filter can fairly quickly detect the sudden power spike (within 15 minutes) and then produce good estimates again. Next, we investigate the filter performance by varying the average error in the attacker’s knowledge of the actual heat recirculation matrix. Specifically, we scale the errors in our default case (20% average error, as shown in Fig. 8) and show the average error, maximum error, and standard deviation in the attacker’s power estimation in Fig. 10(b). We see that if the attacker’s assumed heat recirculation matrix does not deviate too much from the actual one, its power estimation is quite accurate (e.g., only 5% average power estimation error, given 30% average error in the attacker’s knowledge of H_b). The low estimation error is partly because the benign tenants’ power has a large fixed portion, while the attacker only needs to detect temporal variations for timing attacks.

To conclude, despite the attacker’s imperfect observation and process models, the Kalman filter can estimate the benign tenants’ power at runtime reasonably well.

4.2.2 Power Attacks. Next, we present our experimental result on how well the Kalman filter can help the attacker time its attacks.

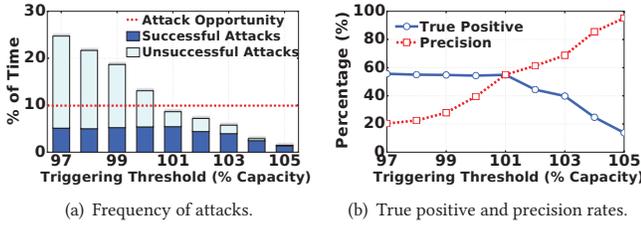


Figure 11: (a) Frequency of power attacks versus the attack triggering threshold. (b) True positive and precision rates versus the attack triggering threshold.

True positive and precision rates. As the attacker cannot launch attacks too frequently (no more than 10% of the time in our default case), true positive and precision rates are important metrics to consider. *True positive rate* is defined as the percentage of available attack opportunities captured by the attacker, while *precision* is the percentage of successful attacks among all the launched attacks. By default, we consider an attack *successful* if the designed power capacity is overloaded for at least 5 minutes.

We first show the frequency of power attacks in Fig. 11(a) by varying the attacker’s triggering threshold. When the attacker sets a lower triggering threshold, it will attack more frequently, detecting more attack opportunities and meanwhile launching more unsuccessful attacks. Thus, as shown in Fig. 11(b), this results in a higher true positive rate but a lower precision rate. To keep the power attacks under 10% of the total time, the attacker can set its triggering threshold at 101% of the designed capacity shared with benign tenants, resulting in a true positive rate of 54% and precision rate of 53%. This represents a significant improvement, compared to the temperature-based attack that only captures 3.9% of the attack opportunities with a precision of 2.1% (Fig. 6(a)).

Impact of T_{attack} and T_{hold} values. The operator’s power contracts vary by data centers, and thus the attacker can adjust its attack strategy parameters (Sec. 3.5). Here, we study the effect of varying T_{attack} and T_{hold} on the attack success rates in Fig. 12. Specifically, we vary one value while keeping the other as default, and set the triggering threshold to launch attacks for no more than 10% of the time. With an increased T_{attack} , the attacker will peak its power for a longer time and intuitively should yield better attack success rates. This holds for $T_{attack} \leq 25$ minutes. However, the true positive and precision rates may decrease as T_{attack} continuously increases, because the attacker may keep on attacking even though the attack opportunity is gone. On the other hand, with an increased T_{hold} , the attacker will wait longer before re-launching an attack, even though an attack opportunity may appear sooner. Hence, we see in Fig. 12(b) that the attack success rates decrease as T_{hold} increases.

Comparison with random attacks. Without a (thermal) side channel, the attacker may launch random attacks, possibly within a narrower time frame if coarse timing is available (Sec. 3.2). Random attacks can also capture a *benign* tenant which unintentionally peaks its power usage. We now compare our timed attack with random attack on a yearly trace in Fig. 13. Intuitively, randomly attacking for $X\%$ of the time should capture $X\%$ of the available

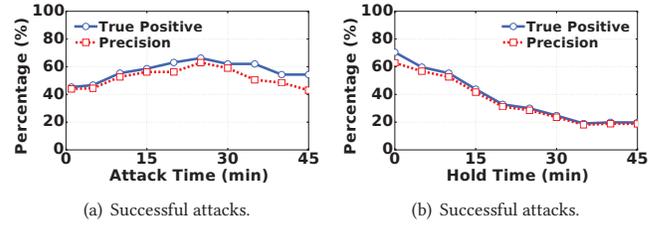


Figure 12: Attack success rates for different timer values.

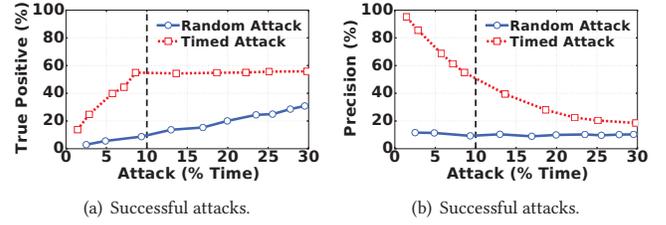


Figure 13: Comparison with random attacks.

attack opportunities, with a fixed precision rate that is the same as the probability of attack opportunities. This can be seen from Fig. 13, where the small disturbances are due to empirical evaluations. Nonetheless, our timed attack significantly outperforms random attacks, especially for limited attacking time less than 10% of the time. This highlights the necessity of a (thermal) side channel as well as the danger of maliciously timed power attacks. Note that after an initial increase, the true positive rate of our timed attacks remains saturated even when the attacker attacks more frequently (which also means a lower precision rate). This is because the total available attack opportunities are the same and some of them can span a relatively longer (e.g., 20 minutes), but we do not allow the attacker to attack consecutively (Sec. 3.5).

Impact of the attacker size. Naturally, a larger attacker with a higher capacity subscription can launch more successful attacks and make the power infrastructure less reliable. In Fig. 14(a), we show the impact of attacker size on the available attack opportunities and its attack success rates. We keep the benign tenants’ total capacity fixed and scale up the attacker’s capacity to the different percentages of total subscribed capacity. We also keep the total attacking time at the default 10%. Naturally, the number of attack opportunities increases with the attacker size, as the attacker can create higher power spikes. We also see that as the attacker has more servers, the true positive rate can go down while the precision increases. This is because, although there are more attack opportunities, the total attacking time remains the same, thus possibly resulting in a lower true positive rate. At the same time, as there are more opportunities, the precision rate goes up.

Next, in Fig 14(b), we show the annual cost impact (following Sec. 2.3) with varying sizes of the attacker. The attacker needs to pay more as rent when its subscribed capacity is larger, with an annual cost of \$48.8k at 5% size up to \$308.9k at 25% size (assuming a capacity leasing cost of \$150/kW/month, energy cost of 10

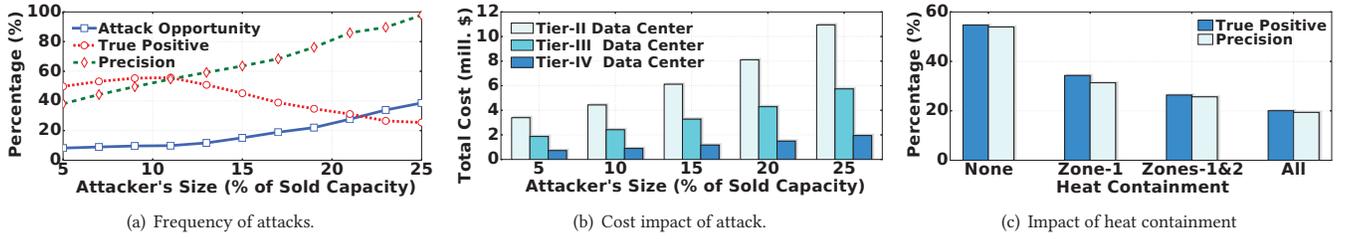


Figure 14: (a) Statistics of attack opportunity and attack success. (b) Expected annual loss due to power attacks incurred by the data center operator and affected tenants (200kW designed capacity oversubscribed by 120%). (c) Even with heat containment, the thermal side channel can still assist the attacker with timing power attacks.

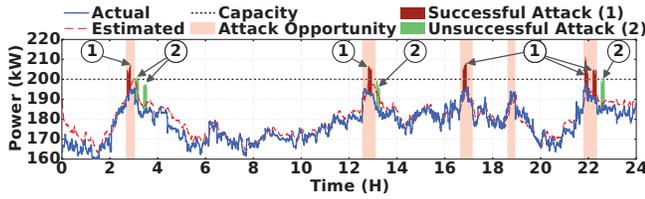


Figure 15: Illustration of different attack scenarios.

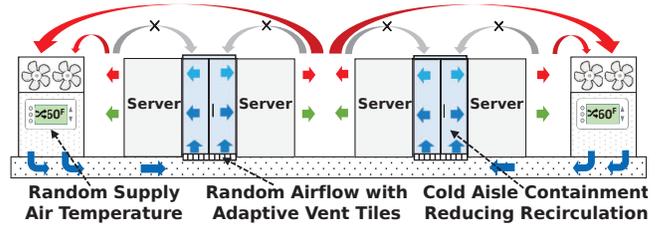


Figure 16: Degrading the thermal side channel.

cents/kWh, and server cost of \$1500 per 250W server amortized over 3 years). But these costs are just a fraction (varying between 1.44% and 15.88% depending on the attacker size and data center tier) of the total cost borne by the operator and affected tenants due to the compromised data center availability. On the other hand, a larger attacker can create more emergencies and cause more damages to the data center. We see that, by spending in the order of 100 thousand dollars per year, the attacker can cost the target data center an annual loss in the order of millions of dollars.

Impact of heat containment. While full heat containment is rare in multi-tenant data centers, it may be partially implemented (see Fig. 4). Here, we study the impact of different degrees of heat containment on the timing accuracy of attacks. We consider three different cases, where one, two and three zones have heat containment, respectively. As heat containment can reduce, but not completely eliminate, heat recirculation [44], we consider that the corresponding heat recirculation impact is reduced by 90% when a zone has heat containment. We see in Fig. 14(c) that heat containment reduces both the true positive rate and precision. Nonetheless, this is still higher than random attacks.

Illustration of different attack scenarios. Finally, we show a snapshot of the power attack trace in Fig. 15 to illustrate what *would* happen had attacks been launched based on the strategy described in Sec. 3.5. We see some successful attacks that can create prolonged overloads of the shared capacity. Note that an actual outage may not always occur after a capacity overload due to infrastructure redundancy, but if it does occur, the power trace will differ after the outage incident. There are also unsuccessful attacks in Fig. 15 due to overestimates of the benign tenants' aggregate power demand, which fails to overload the designed capacity. In addition, there are missed opportunities around the 19-th hour.

5 DEFENSE STRATEGY

Given the danger of power attacks, a natural question follows: *how can a multi-tenant data center operator better secure its power infrastructure against power attacks?* In this section, we discuss a few possible defense strategies.

5.1 Degrading Thermal Side Channel

Since the thermal side channel resulting from heat recirculation is instrumental to time power attacks, the first natural defense strategy would be degrading the side channel. This can make the attacker estimate the benign tenants' power usage with more errors, thus misguiding the attacker's power attacks. Towards this end, the data center can either randomize the cooling system set point or reduce heat recirculation through heat containment.

Randomizing supply air temperature. Supply air temperature $T_{sup}(t)$ is an important parameter for the attacker's observation model in (2), and its randomization might *confuse* the attacker. However, the attacker can easily set $T_{sup}(t)$ as a new state to estimate along with the states of benign tenants' power consumption in the Kalman filter, and estimate it fairly accurately. Thus, randomizing $T_{sup}(t)$ does not offer a good protection against power attacks. Further, it can decrease the cooling efficiency (due to, e.g., unnecessarily low temperature settings).

Randomizing supply airflow. Another approach is to make the actual heat recirculation process more uncertain to the attacker. In particular, randomizing the supply airflow can make the attacker's knowledge of the heat recirculation matrix more erroneous. This requires the data center operator install adaptive vent tiles and carefully adjust their opening without server overheating, incurring a high control complexity [46]. Moreover, as the attacker only

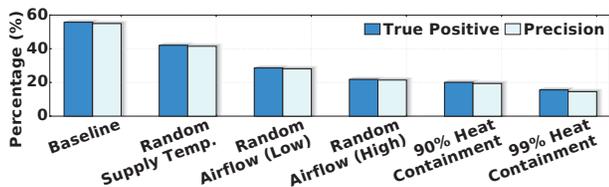


Figure 17: True positive and precision rates of different defense strategies. “Low”/“high” indicates the amount of randomness in supply airflows. “x%” heat containment means x% of the hot air now returns to the CRAH unit directly.

needs to know the benign tenants’ aggregate power rather than individual power, the Kalman filter performance is still reasonably good in the presence of supply airflow randomization, making this strategy only moderately effective.

Heat Containment. While container-based design (e.g., enclosing tenants’ servers in a shipping container) can isolate thermal recirculation across tenants [58], it is costly to implement and rarely used in multi-tenant data centers [59]. Instead, the data center operator typically decreases heat recirculation by sealing the cold or hot aisles [43, 60]. Cold/hot aisle containment has a reasonably low capital expense but, due to tenants’ heterogeneous racks, only has limited adoption in multi-tenant data centers (especially existing ones) as shown in Fig. 4 [24]. Nonetheless, once heat containment is successfully installed, only very little hot air can recirculate and there is no control needed at runtime. Thus, heat containment can be effective with a low capital expense.

We illustrate the aforementioned defense strategies in Fig. 16. We also quantify the effectiveness of different defense strategies by investigating their impacts on the attacker’s true positive and precision rates of successful attacks. The results are shown in Fig. 17, where “Baseline” is the current status quo without our discussed defenses. For all the defenses, the attacker uses the same attack strategy as discussed in Section 3.5. We see that heat containment is the most effective strategy, while randomizing the supply air temperature has little effect in preventing power attacks. In particular, with 99% heat containment (i.e., only 1% hot air recirculates), the attacker’s timing accuracy through the thermal side channel is only marginally better than random attacks.

We recommend *heat containment* as the “best” defense strategy due to its high effectiveness, low cost and zero management at runtime. Thus, besides efficiency [43], securing the power infrastructure against power attacks now becomes another compelling reason for multi-tenant data centers to adopt heat containment.

5.2 Other Countermeasures

There also exist other countermeasures to secure a multi-tenant data center against power attacks. A straightforward approach is to not oversubscribe the power infrastructure, thus eliminating the vulnerabilities and attack opportunities. But, this comes at a significant revenue loss for multi-tenant data center operators and installing extra capacity can be particularly challenging in existing data centers. Another approach is to increase the level of redundancy. Nonetheless, the attacker can still compromise the

long-term *designed* availability, which essentially translates into a capital loss for the operator (Table 1).

It is also important to detect the malicious attacker as early as possible and then evict it. While the power usage illustrated in Fig. 2(c) does not violate the operator’s contract and can be a benign tenant’s power pattern, continuously having such a usage pattern may be suspicious. Concretely, the operator may pay special attention to the high aggregate power periods and closely monitor which tenant has the highest contribution to those periods.

Finally, the operator can take other measures or implement a combination of the above strategies to secure its infrastructure against power attacks. This is an interesting research direction for our future study.

6 RELATED WORK

Power oversubscription is economically compelling but can result in occasional emergencies that require power capping to handle [6]. For example, well-known power capping techniques include throttling CPU frequencies [8, 51], reducing workloads [15], among others. Unfortunately, these approaches cannot be applied by a multi-tenant data center operator due to the operator’s lack of control over tenants’ servers.

There have been many studies on making the *cyber* part of a data center more secure. For example, defending data center networks against DDoS attacks [27, 28] and protecting user privacy against side channel attacks [29, 30, 61, 62] have both received much attention.

In parallel, data center *physical* security has been gaining attention quickly in recent years. For example, [63] studies defending servers against human intrusion and attacks. More recently, [12, 32] attempt to intentionally create power emergencies in an owner-operated data center through VMs. Nonetheless, malicious VM workloads may not all be placed together to create high and prolonged spikes, and the operator can use server power and VM placement control knobs in place to safeguard the power infrastructure [29].

In contrast, we consider a multi-tenant data center where a malicious tenant can subscribe enough power capacity to create extended and severe power emergencies. Further, the data center operator has no control of tenants’ servers and thus, cannot apply power capping to mitigate power attacks. More importantly, unlike [12, 31, 32], we exploit a co-residency thermal side channel resulting from the unique heat recirculation to launch *well-timed* power attacks. Thus, our work represents the first effort to defend *multi-tenant* power infrastructures against power attacks.

Our work also makes contributions to the literature on multi-tenant data center power management. Concretely, the existing studies have all been *efficiency*-driven, such as reducing energy costs [3], increasing power utilization [17] and minimizing social cost for demand response [64]. In contrast, our work focuses on the power infrastructure security, a neglected but very important issue in multi-tenant data centers.

Finally, we discuss if the attacker can alternatively exploit other side channels. In general, when workload increases, the server power also increases and so does the latency [65]. Thus, request response time might be a cyber side channel: a higher response

time might indicate that the tenant is having a higher workload and hence, power usage, too. Nonetheless, many tenants do not even have any services open to the public [65, 66]. Thus, the measured response time contains little, if any, information about the *aggregate* power usage of multiple benign tenants. Further, the attacker might infer the benign tenants’ power usage based on its detected voltage/current changes. However, a multi-tenant data center delivers highly *conditioned* power to tenants’ servers, and the internal wiring topology (e.g., Fig. 1) may not be known to the attacker. In any case, we make the first effort to study power attacks in multi-tenant data centers by exploiting a thermal side channel, which can complement other side channels (if any) and assist the attacker with timing its attacks more accurately.

7 CONCLUDING REMARKS

In this paper, we study a new attack – maliciously timing peak power usage to create emergencies and compromise the availability of a multi-tenant data center. We demonstrate that an oversubscribed multi-tenant data center is highly vulnerable to maliciously timed high power loads. We identify a thermal side channel due to heat recirculation that contains information about the benign tenants’ power usage and design a Kalman filter guiding the attacker to precisely time its attacks for creating power emergencies. Our experiments show that the attacker can capture 54% of all attack opportunities with a precision rate of 53%, highlighting a high success rate and danger of well-timed power attacks.

ACKNOWLEDGMENT

This work was supported in part by the U.S. NSF under grants CNS-1551661, CNS-1565474, ECCS-1610471, AitF-1637598, CNS-1518941, and CNS-1319820.

APPENDIX

A CIRCUIT BREAKER TRIP DELAY CURVE

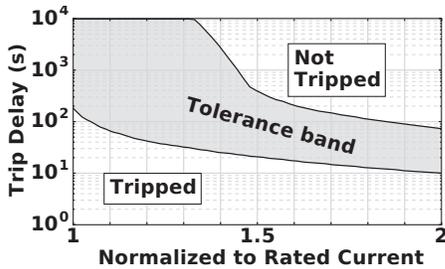


Figure 18: Circuit breaker trip delay [34].

We show in Fig. 18 the circuit breaker trip delay curve according to [34]. It can be seen that although the circuit breaker can tolerate transient overloads, prolonged overloads will make it trip and result in power outages.

B DATA CENTER LAYOUT

We follow the design of HP Labs [46] and show the indoor part of our considered data center space in Fig. 19. The total designed

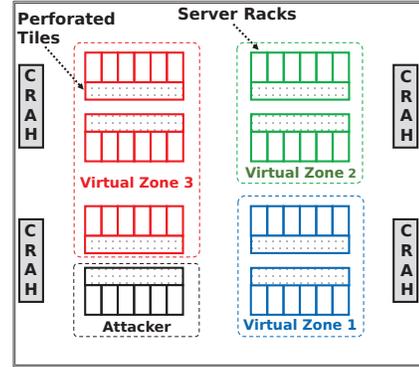


Figure 19: Data center layout.

capacity under consideration is 200kW and, according to the industry average [23], oversubscribed by 120%. We consider the attacker divides the shared data center space into four different virtual zones (three for benign tenants and one for the attacker). Zones 1 and 2 have 12 server racks each. Zone 3 has 18 server racks, while the attacker occupies the 4th zone with 6 racks. Each rack has 20 servers and a power capacity of 5kW. There are four CRAH units that supply cold air to servers through perforated tiles.

C BREAKDOWN OF TEMPERATURE READING

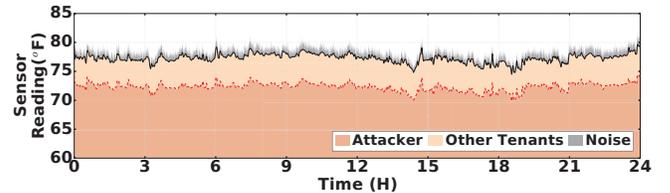


Figure 20: Breakdown of readings at sensor #1 (Fig. 5(a)).

The impact of heat recirculation is illustrated in Fig. 20, where we consider that the CRAH delivers cold air at 60°F through perforated tiles and show the breakdown of temperature readings monitored at sensor #1 through CFD analysis (details in Section 4.1). Note that the breakdown is for demonstrating the impact of benign tenants’ power usage on the attacker’s server inlet temperature, while it is not accurately known to the attacker in practice. We see that the benign tenants’ servers have a noticeable impact on the attacker’s server inlet temperature, potentially leaking the benign tenants’ power usage information to the attacker at runtime.

D SNAPSHOT OF TEMPERATURE-BASED POWER ATTACKS

The attacker launches a 10-minute attack whenever its average temperature reading exceeds 76.3°F for at least 1 minute. The snapshot in Fig. 21 shows that all attacks are unsuccessful.

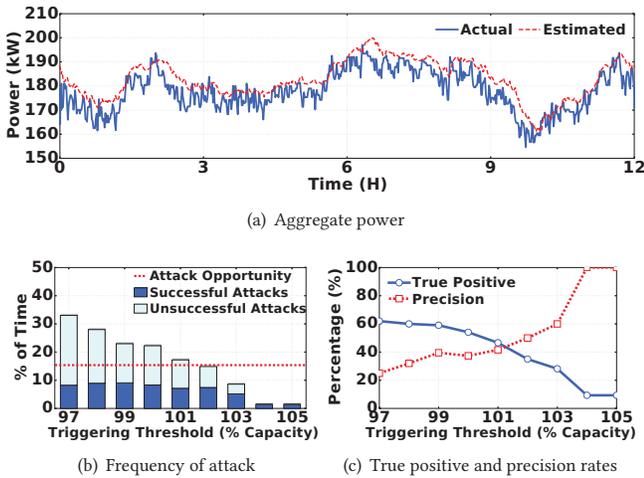


Figure 23: Detection on an alternative power trace.

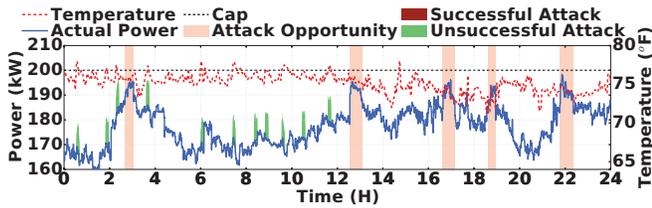


Figure 21: Temperature-based power attack. All attacks are unsuccessful.

E HEAT RECIRCULATION MODEL ASSUMED BY THE ATTACKER

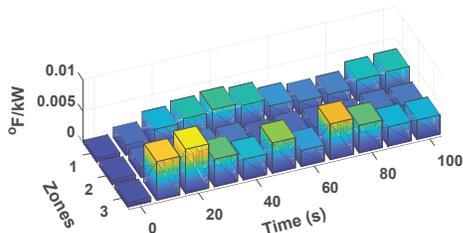


Figure 22: The attacker's heat recirculation model: zone-wise temperature increase at sensor #1 (Fig. 5(a)).

We show the attacker's estimate of zone-based temperature increase impact ($^{\circ}\text{F}/\text{kW}$) at one of its sensors in Fig. 22. The bars show the impact of different zones' power on the sensor reading with time.

F DETECTION STATISTICS FOR ALTERNATE POWER TRACE

We evaluate the performance of Kalman filter and attack success rate using an alternative set of power traces taken from [15, 57].

There are more attack opportunities in the alternate power trace than in the default case. We present a 12-hour snapshot of the actual aggregate power and estimated value in Fig. 23(a). We also show the available attack opportunity, successful attacks and unsuccessful attacks for different triggering thresholds in Fig. 23(b). We see that, even with more attack opportunities, the successful attack is slightly lower than the default case. This is mainly because attack opportunities last longer in our alternate power trace but the attacker restrains itself from launching consecutive attacks to stay stealthy. The corresponding true positive and precision rates are shown in Fig. 23(c), which are comparable to our default case and demonstrate the effectiveness of the thermal side channel in terms of timing power attacks.

REFERENCES

- [1] "Colocation market - worldwide market forecast and analysis (2013 - 2018)," <http://www.marketsandmarkets.com/ResearchInsight/colocation.asp>.
- [2] NRDC, "Scaling up energy efficiency across the data center industry: Evaluating key drivers and barriers," *Issue Paper*, Aug. 2014.
- [3] M. A. Islam, H. Mahmud, S. Ren, and X. Wang, "Paying to save: Reducing cost of colocation data center via rewards," in *HPCA*, 2015.
- [4] Apple, "Environmental responsibility report," 2016.
- [5] Y. Sverdlik, "Silicon valley: A landlord's data center market," in *DataCenterKnowledge*, February 2015.
- [6] X. Fan, W.-D. Weber, and L. A. Barroso, "Power provisioning for a warehouse-sized computer," in *ISCA*, 2007.
- [7] A. Greenberg, J. Hamilton, D. A. Maltz, and P. Patel, "The cost of a cloud: Research problems in data center networks," *SIGCOMM Comput. Commun. Rev.*, vol. 39, Dec. 2008.
- [8] Q. Wu, Q. Deng, L. Ganesh, C.-H. R. Hsu, Y. Jin, S. Kumar, B. Li, J. Meza, and Y. J. Song, "Dynamo: Facebook's data center-wide power management system," in *ISCA*, 2016.
- [9] B. Kleyman, "Knowing when dedicated data centers are just not enough," *Data Center Frontier*, December 2015.
- [10] Leagle.com (Case No. 5:13-cv-03093-PSG), "Layton v. Terremark North America, LLC," June 2014.
- [11] Hornbaker Group, "Determining kilowatt capacity of data center space," <http://www.hornbakergroup.com/pdf/Considerations-when-leasing-Data-Center-space-by-the-kilowatt.pdf>.
- [12] C. Li, Z. Wang, X. Hou, H. Chen, X. Liang, and M. Guo, "Power attack defense: Securing battery-backed data centers," in *ISCA*, 2016.
- [13] S. Govindan, D. Wang, A. Sivasubramaniam, and B. Urgaonkar, "Aggressive datacenter power provisioning with batteries," *ACM Trans. Comput. Syst.*, vol. 31, pp. 2:1–2:31, Feb. 2013.
- [14] D. Wang, C. Ren, and A. Sivasubramaniam, "Virtualizing power distribution in datacenters," in *ISCA*, 2013.
- [15] G. Wang, S. Wang, B. Luo, W. Shi, Y. Zhu, W. Yang, D. Hu, L. Huang, X. Jin, and W. Xu, "Increasing large-scale data center capacity by statistical power control," in *EuroSys*, 2016.
- [16] Ponemon Institute, "2016 cost of data center outages," 2016, <http://goo.gl/6mBFTV>.
- [17] M. A. Islam, X. Ren, S. Ren, A. Wierman, and X. Wang, "A market approach for handling power emergencies in multi-tenant data center," in *HPCA*, 2016.
- [18] CNN, "Delta: 5-hour computer outage cost us \$150 million," Sep. 07 2016 (<http://money.cnn.com/2016/09/07/technology/delta-computer-outage-cost/>).
- [19] Colocation America, "Data center standards (Tiers I-IV)," 2017, <https://www.colocationamerica.com/data-center/tier-standards-overview.htm>.
- [20] Uptime Institute, "Tier certifications," <https://uptimeinstitute.com/TierCertification/>.
- [21] 365DataCenters, "Master services agreement," <http://www.365datacenters.com/master-services-agreement/>.
- [22] Internap, "Colocation services and SLA," <http://www.internap.com/internap/wp-content/uploads/2014/06/Attachment-3-Colocation-Services-SLA.pdf>.
- [23] United States District Court, "Layton v. Terremark North America, LLC," 2014.
- [24] Uptime Institute, "Data center industry survey," 2014.
- [25] M. Jonas, R. R. Gilbert, J. Ferguson, G. Varsamopoulos, and S. K. Gupta, "A transient model for data center thermal prediction," in *IGCC*, 2012.
- [26] J. Moore, J. Chase, P. Ranganathan, and R. Sharma, "Making scheduling 'cool': Temperature-aware workload placement in data centers," in *USENIX ATC*, 2005.
- [27] J. Mirkovic and P. Reiher, "A taxonomy of DDoS attack and ddos defense mechanisms," *SIGCOMM Comput. Commun. Rev.*, vol. 34, pp. 39–53, Apr. 2004.

- [28] S. Yu, Y. Tian, S. Guo, and D. O. Wu, "Can we beat ddos attacks in clouds?," *IEEE Trans. Parallel and Distrib. Syst.*, vol. 25, pp. 2245–2254, September 2014.
- [29] S.-J. Moon, V. Sekar, and M. K. Reiter, "Nomad: Mitigating arbitrary cloud side channels via provider-assisted migration," in *CCS*, 2015.
- [30] Y. Zhang, A. Juels, M. K. Reiter, and T. Ristenpart, "Cross-vm side channels and their use to extract private keys," in *CCS*, 2012.
- [31] X. Gao, Z. Gu, M. Kayaalp, D. Pendarakis, and H. Wang, "ContainerLeaks: Emerging security threats of information leakages in container clouds," in *DSN*, 2017.
- [32] Z. Xu, H. Wang, Z. Xu, and X. Wang, "Power attack: An increasing threat to data centers," in *NDSS*, 2014.
- [33] Telecommunications Industry Association, "Data center standards overview," *TIA 942*, 2005 (amended in 2014).
- [34] Raritan, "Data center power overload protection," *White Paper*, 2016.
- [35] W. P. Turner, J. H. Seader, and K. G. Brill, "Tier classifications define site infrastructure performance," *Uptime Institute White Paper 17*, 2006.
- [36] N. Rasmussen, "Overload protection in a dual-corded data center environment," *APC White Paper 206*, 2014.
- [37] Reuters, "British Airways \$100M outage was caused by worker pulling wrong plug," Jun. 02 2017.
- [38] Q. Pu, G. Ananthanarayanan, P. Bodik, S. Kandula, A. Akella, P. Bahl, and I. Stoica, "Low latency geo-distributed data analytics," in *SIGCOMM*, 2015.
- [39] Z. Wu, M. Butkiewicz, D. Perkins, E. Katz-Bassett, and H. V. Madhyastha, "Spanstore: Cost-effective geo-replicated storage spanning multiple cloud services," in *SOSP*, 2013.
- [40] CBRE, "Q4 2013: National data center market update," 2013.
- [41] Autodesk, "CFD simulations of data centers," <http://auworkshop.autodesk.com/library/cfd-aec/cfd-simulations-data-centers>.
- [42] R. A. Steinbrecher and R. Schmidt, "Data center environments: Ashrae's evolving thermal guidelines," *ASHRAE Technical Feature*, pp. 42–49, December 2011.
- [43] D. L. Moss, "Dynamic control optimizes facility airflow delivery," *Dell White Paper*, March 2012.
- [44] S. V. Patankar, "Airflow and cooling in a data center," *Journal of Heat Transfer*, vol. 132, p. 073001, July 2010.
- [45] T. Evans, "Fundamental principles of air conditioners for information technology," *Schneider Electric White Paper 57*, March 2015.
- [46] Z. Wang, A. McReynolds, C. Felix, C. Bash, C. Hoover, M. Beitelmal, and R. Shih, "Kratos: Automated management of cooling capacity in data centers with adaptive vent tiles," in *ASME International Mechanical Engineering Congress and Exposition*, 2009.
- [47] D. Kennedy, "Optimizing capacity and efficiency in a diverse and variable load environment," *TATE*, August 2010.
- [48] N. Rasmussen, "Guidelines for specification of data center power density," *Schneider Electric White Paper 120*, April 2015.
- [49] R. McFarlane, "Controversial hot aisle containment practices," *Techtarget*, June 2014.
- [50] J. Novet, "Colocation providers, customers trade tips on energy savings," Nov. 2013, <http://www.datacenterknowledge.com/>.
- [51] L. Li, W. Zheng, X. D. Wang, and X. Wang, "Coordinating liquid and free air cooling with workload allocation for data center power minimization," in *ICAC*, 2014.
- [52] Autodesk, "Computational fluid dynamics," <http://www.autodesk.com/products/cfd/overview>.
- [53] Q. Tang, S. K. S. Gupta, and G. Varsamopoulos, "Energy-efficient thermal-aware task scheduling for homogeneous high-performance computing data centers: A cyber-physical approach," *IEEE Trans. Parallel Distrib. Syst.*, vol. 19, pp. 1458–1472, Nov. 2008.
- [54] S. S. Haykin, *Kalman Filtering and Neural Networks*. New York, NY, USA: John Wiley & Sons, Inc., 2001.
- [55] Google, "Google's Data Center Efficiency," <http://www.google.com/about/datacenters/>.
- [56] D. G. Feitelson, D. Tsafir, and D. Krakov, "Experience with using the parallel workloads archive," *Journal of Parallel and Distributed Computing*, vol. 74, no. 10, pp. 2967–2982, 2014.
- [57] Parallel Workloads Archive, <http://www.cs.huji.ac.il/labs/parallel/workload/>.
- [58] A. Qouneh, C. Li, and T. Li, "A quantitative analysis of cooling power in container-based data centers," in *IISWC*, 2011.
- [59] Y. Sverdlik, "Microsoft moves away from data center containers," in *DataCenter Knowledge*, April 2016.
- [60] T. Evans, "The different technologies for cooling data centers," http://www.apcmmedia.com/salestools/VAVR-5UDTU5/VAVR-5UDTU5_R2_EN.pdf.
- [61] M. Guri, M. Monitz, Y. Mirski, and Y. Elovici, "Bitwhisper: Covert signaling channel between air-gapped computers using thermal manipulations," in *IEEE Computer Security Foundations Symposium*, 2015.
- [62] P. C. Kocher, J. Jaffe, and B. Jun, "Differential power analysis," in *CRYPTO*, 1999.
- [63] J. Szefer, P. Jankhedkar, Y.-Y. Chen, and R. B. Lee, "Physical attack protection with human-secure virtualization in data centers," in *Dependable Systems and Networks Workshops (DSN-W)*, 2012.
- [64] N. Chen, X. Ren, S. Ren, and A. Wierman, "Greening multi-tenant data center demand response," in *IFIP Performance*, 2015.
- [65] L. A. Barroso, J. Clidaras, and U. Hoelzle, *The Datacenter as a Computer: An Introduction to the Design of Warehouse-Scale Machines*. Morgan & Claypool, 2013.
- [66] J. dePreaux, "Wholesale and retail data centers - North America and Europe - 2013," *IHS*, Jul. 2013, <https://technology.ihs.com/api/binary/492570>.
- [67] 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," in *ISCA*, 2015.