

A First Look at Power Attacks in Multi-Tenant Data Centers

[Extended Abstract]

Mohammad A. Islam
University of California,
Riverside

Shaolei Ren
University of California,
Riverside

Adam Wierman
California Institute of
Technology

ABSTRACT

Oversubscription increases the utilization of expensive power infrastructure in multi-tenant data centers, but it can create dangerous emergencies and outages if the designed power capacity is exceeded. Despite the safeguards in place today to prevent power outages, this extended abstract demonstrates that multi-tenant data centers are vulnerable to well-timed power attacks launched by a malicious tenant (i.e., attacker). Further, we show 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. We develop a state-augmented Kalman filter that guides an attacker to precisely time its power attacks at moments that coincide with the benign tenants’ high power demand, thus overloading the designed power capacity. Our experimental results show that an attacker can capture 53% of all attack opportunities, significantly compromising the data center availability.

1. INTRODUCTION

Multi-tenant data centers are shared data center facilities where tenants house their physical servers while the data center operator is responsible for managing the non-IT support systems such as power distribution and cooling. They are a very important segment of data centers and adopted widely by many industry sectors. For example, Apple has 25% of its servers in multi-tenant data centers. As of today, there are nearly 2,000 multi-tenant data centers in the U.S., consuming five times the energy of Google-type data centers altogether.

To accommodate the fast growing Internet and cloud-based services, multi-tenant data center operators are under a constant pressure to expand and/or build new facilities. However, data center infrastructure is very expensive to build because of its high availability requirement, costing around US\$10 ~ 25 per watt of IT critical power delivered (cooling power is separate from IT power) and taking up more than 1.5 times of the total electricity cost over its lifespan. In addition, long-time-to-market and local power grid constraints also pose significant hurdles for increasing capacity in multi-tenant data centers.

Consequently, to maximize the utilization of expensive data center infrastructures, multi-tenant data center operators commonly oversubscribe the power infrastructure by selling capacity to more tenants than can be supported. Even owner-operated data centers, such as Facebook, oversubscribe power infrastructures to defer/reduce the need of building new

capacities. The industry standard oversubscription ratio is 120%, and recent studies have begun to suggest even more aggressive oversubscription [3].

Power oversubscription increases capacity utilization and significantly reduces capital expenses. But, it comes with a dangerous consequence of overloading the designed capacity (a.k.a., *power emergencies*) when the power demand of multiple tenants peaks simultaneously. In fact, even short-term overloads over a few minutes can lead to tripped circuit breakers and costly data center outages (e.g., Delta Airline’s data center power outage resulted in a US\$150 million loss [1]).

Power infrastructure redundancy is very common to ensure a high availability in today’s data centers. While it can safeguard the data center against power emergencies by taking over some overloads, redundancy protection is lost during an emergency. For example, if with an emergency, a fully-redundant Tier-IV data center can experience an outage when either the primary or redundant infrastructure fails; otherwise, an outage occurs only when *both* the primary and redundant infrastructures fail. The loss of infrastructure redundancy can significantly increase the outage risk of a fully redundant Tier-IV data center by 280+ times [2]. Even though emergencies only occur for 5% of the time, the availability of a Tier-IV data center can be downgraded to that of a Tier-II data center, which means a nearly 50% capital loss for the data center operator.

The severe consequences of power emergencies have led to the development of various power capping techniques such as CPU throttling. Nonetheless, the lack of control over tenants’ servers renders such techniques inapplicable in multi-tenant data centers. Concretely, a power emergency may still occur, even though all the tenants are using power within their purchased capacities due to the operator’s oversubscription decision. Multi-tenant data center operators, therefore, have taken other precautions by imposing restrictions on each tenant’s “normal” power usage to be below a certain percentage (e.g., 80%) of its subscribed capacity, which limits frequent and/or constant usage of a tenant’s full capacity. This contractual constraint effectively reduces the probability of simultaneous peaks of tenants’ power usage, keeping the risk of power emergencies at a very low level. Hence, despite oversubscription, power infrastructure in multi-tenant data centers is considered *safe*.

Contributions. This extended abstract summarizes our recent work [4]. *Our goal is to highlight that, despite the safeguards in place today, multi-tenant data centers are vulnerable to well-timed malicious power attacks that can cause a huge financial loss for the data center operator as well as affected tenants.* More specifically, a malicious tenant (i.e., attacker), which can be the competitor of the target data center and does not run any useful workloads, can



Figure 1: Illustration of opportunity and power attack.

intentionally increase power to its peak/subscribed capacity in an attempt to create power emergencies, when it detects an already high utilization of the shared capacity. Importantly, the total cost incurred by the attacker, such as server and power capacity costs, is only small fraction (between 1.44% and 15.88%, as shown by [4]) of the total financial loss borne by the operator and benign tenants.

To illustrate this point, we show in Fig. 1 a 24-hour power trace by four representative tenants. These tenants run web and data analysis workloads. The total capacity is 200kW and sold as 240kW following industry standard oversubscription of 120%. The attacker subscribes to 30kW and increases its power to full capacity for 10 minutes whenever the aggregate power approaches the capacity limit. Consequently, we see multiple power emergencies, while the attacker only occasionally peaks its power and meets its contractual constraint.

While power attacks are dangerous, attack opportunities only exist intermittently due to the fluctuation of benign tenants’ aggregate power usage, as illustrated in Fig. 1. Thus, a key question is: *how does the attacker detect an attack opportunity?*

Keeping a constant high power will capture all attack opportunities, but it is not allowed by operator and can lead to the attacker’s eviction. Because of intermittency of attack opportunities, randomly attacking is not likely to be successful either. Further, even though some coarse windows of attack opportunities (e.g., possibly peak traffic hours) can be identified and help the attacker locate the attack opportunities within a smaller time frame, the actual attack opportunity is short-duration and may not last throughout the entire coarse window. Thus, the attacker needs to *precisely* time its attacks to coincide with other benign tenants’ high power usage. Nonetheless, this may seem impossible, as the attacker does not have access to the operator’s power meters to monitor the benign tenants’ power usage at runtime.

We observe that the physical co-location of the attacker and benign tenants in shared data center spaces results in a prominent thermal side channel which carries benign tenants’ power usage information. Concretely, almost all the power consumed by a server is converted into heat, and due to the lack of complete heat containment in many data centers (see Fig. 3), some of the hot air can travel to other servers (a.k.a. *heat recirculation*) and increases their inlet temperatures. Thus, an attacker can easily conceal temperature sensors in its servers to monitor the inlet temperatures, extracting information of benign tenants’ runtime power usage.

A naive strategy for the attacker is to look at server inlet temperature and launch power attacks whenever the inlet temperature is sufficiently high. But, even with the same power consumption, a server closer to the attacker can cause a greater temperature increase at the attacker’s server inlet than a server that is farther away. Hence, a high temperature reading at the attacker’s inlet does not necessarily indicate a high aggregate power consumption of benign tenants.

We demonstrate that, by leveraging the knowledge of the layout of the target data center (through a maintenance visit of its own

servers) and computational fluid dynamics (CFD), the attacker can obtain a rough idea of the heat recirculation process and use a *state-augmented* Kalman filter to extract the hidden information about benign tenants’ power usage contained in the thermal side channel. Although the attacker’s CFD analysis only provides limited and imprecise knowledge of the actual heat recirculation process, our experiments show that the thermal side channel can assist the attacker with successfully capturing 54% of all the attack opportunities with a precision rate of 53%, significantly threatening the data center availability.

It is also important to note that there might also exist other side channels. For example, a high response time of benign tenants’ services may indicate a high server utilization and power usage, but response time is also affected by multiple factors irrelevant of power and many tenants do not even have any user-facing services for the attacker to exploit. Further, a data center has a complex internal wiring topology (e.g., “wrapped” for N+1 redundancy) that is unknown to the attacker, and hence inferring benign tenants’ power usage from the shared data center power distribution system can be challenging. In any case, we make the first effort to exploit a side channel — thermal side channel, which can complement other side channels (if any) and assist the attacker in timing its attack more accurately.

In conclusion, the *key novelty* of our work is that it is the first study to consider an *adversarial* setting in multi-tenant data centers— well-timed power attacks by exploiting a thermal side channel. There are a small but quickly expanding set of papers [5, 7] that attempt to create malicious virtual machines to overload the power capacity in an owner-operated data center. In sharp contrast, our work exploits a unique co-residency thermal side channel to launch *well-timed* power attacks in a multi-tenant data center where an attacker controls physical servers and can easily inject a high power load to create severe power emergencies.

2. EXPLOITING A THERMAL SIDE CHANNEL

In this section, we exploit a thermal side channel to guide the attacker to time its attacks against the shared power infrastructure, significantly compromising the data center availability.

2.1 Threat Model

We consider that an attacker shares the power infrastructure capacity C with other benign tenants. The attacker can increase its power to its maximum subscription capacity by running CPU-intensive workloads. We consider an attack successful if $p_a + p_b > C$ for at least L minutes, where p_a and p_b are the attacker’s and benign tenants’ power. In our evaluation, we use $L = 5$ minutes, which is sufficient to trip a circuit breaker. The attacker can be a competitor of the target multi-tenant data center and wants to cause a million-dollar loss due to compromised data center availability by spending only a small fraction in its capacity subscription and server costs. Note that we do not consider hiding advanced weapons or physically tampering with the power infrastructures for attacks.

2.2 A Thermal Side Channel

We observe a thermal side channel due to heat recirculation in multi-tenant data centers. As illustrated in Fig. 2, inside a typical server room with the raised-floor design, cold air goes into servers through their inlets, absorbs the server heat, and then returns through hot isles to the computer room air handler (CRAH) to be

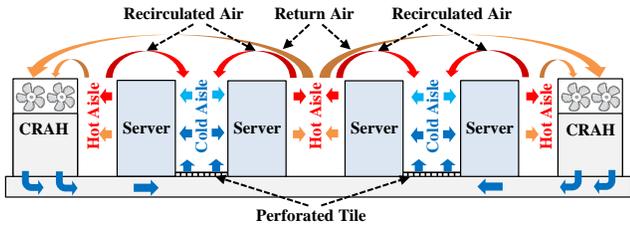


Figure 2: Cooling system overview.

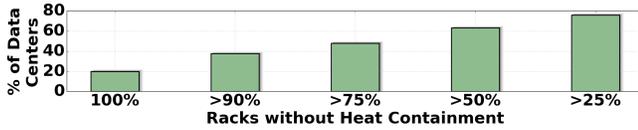


Figure 3: Adoption of heat containment [6].

cooled again. While hot/cold aisle containment can be utilized to guide the airflow and restrict mixing of hot and cold air, it requires homogeneous rack configuration, which is not applicable for many multi-tenant data centers. In addition, heat containment can introduce fire hazards/risks, which many tenants are not willing to undertake. As a result, an open-flow cooling system is widely used in multi-tenant data centers. In fact, according to a survey done by Uptime Institute on 1000+ data centers [6] and shown in Fig. 3, almost 80% of the data centers have at least 25% of their racks without any heat containment and 20% of the data centers have no heat containment at all. Without heat containment, some hot air can travel a few meters to other servers and increases their inlet temperatures, constituting a thermal side channel that conveys some (noisy) information of benign tenants' power usage to the attacker.

2.3 Estimating Benign Tenants' Power

Because of the non-uniform impact of different server racks, a high server inlet temperature does not mean that the benign tenants' power usage is also high. Thus, a model for the heat recirculation process is crucial to extract the benign tenants' power usage.

As the attacker does not know all the details of the data center, having a detailed server-level heat recirculation model is nearly impossible. Thus, we create a *zone*-based linear heat-recirculation model for the attacker based on the widely-used CFD analysis, by considering all servers in a zone have a uniform impact on the attacker's sensors. This zonal consideration significantly reduces the complexity of modeling heat recirculation, but naturally comes at the cost of inaccuracy. Nonetheless, the zone-based model suffices to detect attack opportunities.

Based on a zone-level model, we develop a Kalman filter to estimate benign tenants' runtime power usage, which is hidden in the thermal side channel. Although the attacker's zone-level model only provides a limited view of heat recirculation and can deviate from the actual process, our experiments show that the attacker can still estimate the benign tenants' aggregate power with a high accuracy (e.g., only 3% error on average). This is partly because the attacker only needs to track the benign tenants' power variations and know the *aggregate* value.

2.4 Attack Strategy

To launch its attacks, the attacker sets a triggering threshold on its estimate of aggregate power. The attacker waits for T_{wait} minutes to see if the estimate remains high to avoid attacking during a transient power spike. When the estimate surpasses the

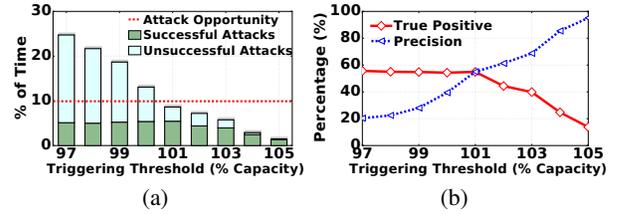


Figure 4: (a) Frequency of power attacks. (b) True positive and precision rates.

triggering threshold for T_{wait} , the attacker starts attacks and keeps its power high for a predetermined time of T_{attack} minutes. In addition, to comply with the contract, the attacker does not re-attack by cooling down for at least T_{hold} minutes, even though it may detect a consecutive attack opportunity.

2.5 Experimental Evaluation

We conduct a CFD analysis to simulate heat recirculation processes in a multi-tenant data center and show the summary of power attacks in Fig. 4(a). We set $T_{attack} = 10$, $T_{wait} = 1$ and $T_{hold} = 10$ minutes in our evaluation. With a lower triggering threshold, the attacker will attack more frequently, detecting more attack opportunities and meanwhile launching more unsuccessful attacks. Thus, as shown in Fig. 4(b), this results in a higher true positive rate (percentage of attack opportunities captured), but a lower precision rate (percentage of successful attacks among all the launched attacks). To keep power attacks under 10% of the total time, the attacker can set its triggering threshold at 101%, resulting in a true positive rate of 54% and a precision rate of 53%.

We only include our key findings in this extended abstract, while details on CFD, Kalman filter, and other results are available in [4].

Acknowledgement

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

3. REFERENCES

- [1] CNN. Delta: 5-hour computer outage cost us \$150 million. Sep. 07 2016.
- [2] Colocation America. Data center standards (Tiers I-IV). 2017.
- [3] S. Govindan, D. Wang, A. Sivasubramaniam, and B. Urgaonkar. Aggressive datacenter power provisioning with batteries. *ACM Trans. Comput. Syst.*, 31(1):2:1–2:31, Feb. 2013.
- [4] M. A. Islam, S. Ren, and A. Wierman. Timing power attacks in multi-tenant data centers. *Tech. Report*, 2017.
- [5] C. Li, Z. Wang, X. Hou, H. Chen, X. Liang, and M. Guo. Power attack defense: Securing battery-backed data centers. In *ISCA*, 2016.
- [6] Uptime Institute. Data center industry survey, 2014.
- [7] Z. Xu, H. Wang, Z. Xu, and X. Wang. Power attack: An increasing threat to data centers. In *NDSS*, 2014.