Establishing the Guidelines for Using Refurbished Hardware in Creating New Data Centers to Lower the Amount of E-Waste

J. Redžepagić, V. Dakić, Z. Morić,
Algebra University College, Zagreb, Croatia
jasmin.redzepagic@algebra.hr, dakic.vedran@algebra.hr, zlatan.moric@algebra.hr,

***Abstract* - This paper presents a study on the guidelines for using refurbished hardware in creating new data centers to lower the amount of electronic waste (e-waste). The research examines the potential of refurbished hardware to provide a cost-effective and eco-friendly alternative to new hardware in data center construction. The study also evaluates the technical capabilities of refurbished hardware and how these factors can affect the performance and security of data centers. Additionally, the paper examines the various refurbishment standaorganizationsrds and certifications available and their suitability for data center use. The study also explores the implications of regulatory compliance and industry standards on the use of refurbished hardware. One of the main goals is to define guidelines on what hardware is viable as reusable in data center environments, especially given that most modern software infrastructure is designed with hardware failures in mind and supports hot-swapping parts of the hardware layer. The findings of this paper intend to create a guide for organisations looking to lower their environmental impact and reduce costs while building new data centers.**

Keywords – e-waste management, reuse, extending server lifecycle, reusing hardware, resilient system

I. INTRODUCTION

The typical design of data centers assumes that all the equipment, including infrastructure, servers, and network equipment, is brand new and unused. Green technologies and ecologically friendly designs for data centers aim to reduce carbon emissions and energy usage, but they don't typically consider the possibility of recycling equipment. This paper argues that used equipment should be included in data center design and explores cases in which reusing previously used hardware is feasible.

II. DATA CENTERS

The data center usually denotes the entire infrastructure that enables the Internet today. According to a few surveys, the number of data centers globally is estimated at around 8000, growing at 5% year-to-year [1],[5].

A typical perspective on data centers is that they consist of computers and other necessary infrastructure to perform a specific task or process. In some instances, specifically configured equipment may be used for particular case scenarios, but for the most part, data centers are initially designed as generic units. The primary focus during data center design is maximizing availability, which is achieved by ensuring the equipment is reliable and easily replaceable. This is accomplished in design; data centers are designed from the ground up to be resilient to failures in every conceivable way.

Design starts with choosing a site for the data center. The usual aim is to ensure that redundant power and connections to the rest of the network exist on the site before we begin building – redundant connections are expensive, especially power supply.

The process of designing a data center then tries to cover other essential things: power consumption needs to be in line with what can be provided by the existing network on site. If more power is needed, it may be required to implement additional power lines.

Cooling is also a significant consideration since it has to be able to control the entire heat output of our equipment, and by design, it also needs to be fault tolerant in the same way as the rest of the equipment. Cooling creates additional requirements in terms of space required for the cooling system and the data center structure since cooling equipment tends to be heavy.

Power and cooling are outside our scope of interest, mainly because there is almost no secondary market since the equipment exploatation period is much longer than for the servers.

The traditional design also involves installing and provisioning networking, storage, and computational infrastructure, which up until recently, was always purchased new. The main reason was failures which complicate data center operations. While this viewpoint is understandable, it has evolved over the past few years.

The main thing that has changed is that the software infrastructure running in the data centers is much more resilient to hardware failure than it was a couple of generations ago. In reality, this means that the hardware running in the data centers right now can experience even significant failures while the data center itself can and will simultaneously be completely operational.

While this may seem normal since resistance to hardware failure is ingrained in the fundamentals of the data center design, it has dramatically changed over time. At first, failures were dealt with as exceptions, and failures rarely happened. This led to the creation of infrastructure designed with a small number of points of failure. This meant a backup system was in place for each system, but the general idea was that double failures were reasonably rare.

A. *Fault Resiliency*

Standards for the data centers specify levels of fault resiliency. For the most part, smaller data centers consider redundancy only in critical systems. Usually, this redundancy is noted as “n+1”, meaning the system can resist a single failure before it ceases to function.

With the rising complexity of today's systems, this had to be changed since the sheer number of installed devices meant that failures became commonplace. This means that instead of functioning with zero failures, almost all of today's systems spend a lot of time in a degraded state. This is a concept that has been discussed for years[2].

This degradation can be seen from the user's perspective as either transparent or detrimental.

The term "transparent degradation" refers to a situation where some parts of a system are not functioning as they should. Still, the user is unaware of any issues and believes the system functions normally. The system can detect that specific infrastructure components are not working correctly, but can work around the problems, so the user is not impacted.

On the other hand, "detrimental degradation" refers to a situation where the user experiences problems using the system, either partially or entirely. This may involve the system being completely offline, or performance being so poor that users cannot complete their tasks. In this scenario, the user is aware of the issue, which is considered a problem.

When building a resilient system, the goal is to create one that can continue functioning despite infrastructure failures. By recognizing that failures are a normal part of the system's operation, rather than exceptions, we can design systems that can withstand more failures. We approach hardware and infrastructure failures as common occurrences rather than exceptions and design our data centers accordingly.

B. *Replacing equipment*

According to the surveys published by the Uptime Institute, servers in today’s data centers are getting older:

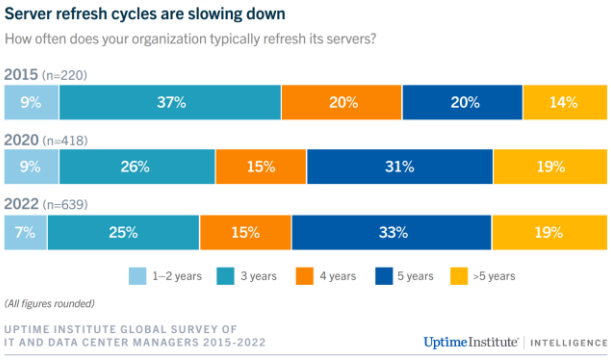

Figure 1 - Server refresh cycles in the last 7 years [1]

Data centers today are less inclined to buy new equipment but also to put older equipment out of commsion. Part of that has to do with “green” initiatives [3], and part is simple economic logic.

What are the reasons the equipment in the data center may be replaced in the first place? One reason for this is when it stops working, which is inevitable when dealing with a large number of units of any hardware. It doesn't matter if we are working with servers, disks, networking equipment, cables, or whatever; the large number of units we will use means that they will be failures and those failures will require replacement.

The second reason for replacing equipment is trying to optimize the process. This usually means that we are trying to either speed up the operations of the data center for a given power envelope or that we are trying to lower the power requirements or the power of the data center while keeping the operation speed at the same level[4]. By operation speed, we usually mean the number of operations or services the given data center can perform when operating normally. In this case, we usually call these replacements "upgrades," but in the end, we are swapping "old" for "new" equipment.

In addition to the other reasons for replacing equipment, a third one involves warranty considerations. When the standard warranty period for a particular unit has expired, repairing or replacing the equipment can become challenging if it fails suddenly.

This is why the equipment that runs out of warranty is usually removed from operation and either disposed of or sold. This is done primarily to ensure that the data center meets its operational agreements with its users. Using equipment without clear service support can create significant logistical issues that may call these agreements into question.

C. *Service levels*

Considering all of the aforementioned points, it is becoming increasingly common for IT infrastructures to require different levels of equipment availability. This is due to the complexity of many systems being used, which often necessitates multiple environments within a data center. One such environment is the production environment for clients, while one or more test environments with less strict availability requirements are also typically used. Even data center users are willing to pay a reduced cost for a test environment with lower availability if it is exclusively used for testing.

By recognizing the varying levels of availability required for different equipment in a data center, it may be beneficial to consider dividing the data center into two separate units. One unit would be dedicated to production, with a high level of availability, while the other would be used for testing and have a lower guaranteed level of availability. This proposal acknowledges the different needs of production and testing environments and aims to optimize the data center's overall efficiency and functionality.

In essence, this means that we will create something reliable but much more ecologically and power friendly with the idea of reusing the hardware that would otherwise be thrown out or not used in a data center environment.

D. Complexity requirements

Analyzing production environments becomes even more complicated when we examine them in depth. Some production environments can be easily replicated without any backups. In contrast, others require lengthy restoration procedures that can take hours, days, or weeks to bring the system back online after a complex failure.

We are going to deal with two examples:

The first one is the usual business-oriented production environment for the company that contains all the systems and services required to run financial systems, CRM (Customer Relationship Manager) systems, front and back-end web systems, and everything else the typical company has. All of this infrastructure needs to be online 24/7.

The other production environment is intended to serve students in their needs for testing and learning environments. Compared to the other environment this student setup is considered production but can be recreated from scratch with only a minimal set of data since all the data in the system is considered test data from the start.

The first environment requires regular backups and a lot of stored and secured data. This means we should design an environment to minimize downtime and avoid restore procedures since they have to move significant amounts of data around.

The second environment can be recreated from scratch in hours using only scripts.

E. System design

One approach that we can take in light of these considerations is to design two different systems. The first would be prepared to run on brand-new servers, while the second would utilize older servers to reduce the cost of acquiring and operating hardware. Typically, older servers are easier and cheaper to acquire and maintain, with more affordable spare parts readily available for upgrades. The price difference between new and old equipment can be significant, often more than enough to compensate for the potential risk of the equipment not performing as expected. If any server fails, it can be easily replaced with another used server at a fraction of the cost.

Until now we have only dealt with the hardware part of the equation since it is the most obvious. In our considerations, we must also take into account software and licensing. While designing a data center, this part of the design is crucial since it has a substantial cost associated with it while at the same time being almost entirely dependent on the hardware design.

Modern software licensing practices are evolving to become more closely tied to the hardware resources a given application utilizes. One notable example is virtualization environments. Licensing is typically determined by factors such as the number of CPU (Central Processing Unit) cores or the total amount of memory allocated for running the environment. This trend has emerged as modern servers continue to expand their capabilities, offering a significantly larger number of CPU cores and increased memory capacity to users.

As licensing fees are often levied on a per-active server basis, encompassing both production and standby servers, this can result in substantial costs for organizations. This financial burden is further worsened when the licensed software is run on refurbished hardware, which is often more affordable than new equipment. In such cases, the total cost of software licensing may even exceed the actual cost of the hardware on which the software is running.

Organizations can better anticipate and manage their expenses by understanding the increasing reliance on hardware resources for software licensing, ensuring a more cost-effective approach to deploying and maintaining software in both production and standby environments.

III. REFURBISHED HARDWARE

When dealing with your refurbished hardware, we must define some things. Other than its positive and negative sides we also need to define the sources we will get it from since we are dealing with the expanding market currently estimated at $5 billion annually. At the same time, the market is almost wholly avoiding equipment manufacturers and relying instead on proxies and second-hand dealers.

A. Sources

There are a couple of ways of acquiring refurbished hardware. The simplest is using one of the many dealers specializing in this business. In the European market, a couple of multimillion-euro companies deal exclusively with refurbished hardware, even offering warranties and spare parts. These companies usually have agreements with big data centers and other large users of enterprise-level equipment.

They usually take care of all the equipment that is either considered broken or old and provide a service where companies can sell their old equipment to the dealer instead of paying for the equipment to be destroyed.

Dealers then evaluate this equipment, check its condition, test it and then usually offer it on the market for a price that varies based on different criteria. Most commonly the most significant criterion is the overall state of the devices.

In many cases, if we are prepared to accept the devices that have scratches and minor dents but are in complete working order, we can acquire equipment for a price that can be measured in a single-digit percentage of the price of the new equipment.

The secondary market also consists of smaller dealers more focused on moving the boxes between companies and end users offering lower prices but providing no additional service and no testing.

There is also the third way of acquiring refurbished hardware by directly contacting companies and data centers, but this is out of the scope of this paper since it

requires lots of additional logistics and provides many additional problems.
Some data centers also rely on their inventory of used equipment since they can stretch the life cycle of the equipment they already have.

B. Reasons to buy

The primary reason for considering refurbished equipment is typically its lower cost, which can be orders of magnitude less than new units, depending on the age of the acquired equipment. Additionally, depending on the use case, refurbished equipment can offer hardware much more capable than new units at the same price point. This is especially applicable to network equipment in smaller environments. This provides the opportunity to purchase more hardware within a set budget or acquire equipment that would otherwise be unattainable at a given price point.

C. Reasons not to buy

Primarily, two things make refurbished equipment non attractive. The obvious one is that we are dealing with the equipment already used, which is pretty unpredictable when it comes to failures. This is a problem that can be dealt with by using systems able to withstand a lot of failures, as we described earlier.
The second issue is considerably more intricate because it involves outdated equipment which, depending on its type and the advancements in this specific market segment, may have a lower cost-to-performance ratio compared to newer equipment. Although we pay less for the older equipment, we receive considerably less performance than we would if we invested in new servers.
Measuring or determining what constitutes “better” equipment is not simple since performance involves different metrics. For example, we may be getting less CPU power but at the same time we may be buying a device that has much more memory and is much more suitable for the task at hand.

IV. Designing a refurbished data center

If we create a data center that relies on refurbished equipment, the primary focus should be the usage case. Typically, when designing a data center, the aim is to cover as many general cases as possible to reach the widest market. However, when working with refurbished hardware, we need to consider the likelihood of hardware failures from the outset and strive to target users or usage cases that can cope with such service disruptions. Therefore, the first step in designing the data center should involve understanding the user and their requirements for the desired service level. Additionally, licenses may pose a challenge since we may need to purchase more resources than necessary to account for the different failure modes of the equipment, and licenses themselves can be expensive since they are directly correlated to the amount of resources we use.

A. Software

Deploying to bare metal servers is almost entirely out of the question when designing with refurbished hardware, so we will probably go with one of the virtualization environments. This is going to be one of the biggest hurdles to overcome since we will pay for the licenses used to cover a lot of resources that will just be spares and hot standbys. Reason for this is that in a regular data center overprovisioning is considered waste of resources and is usually kept to a minimum, so the licensing usually has no provisions for services that are in standby and not in active use.
Our design must do the opposite and consider a lot of overprovisioned hardware to ensure that unexpected hardware failures can be replaced on a live system. That may mean buying licenses to cover the unused hardware simply to have it available if needed.
Aside from virtualization software, other licenses will be the same as in a regular data center.

B. Hardware

Once initial user requirement specifications are defined we should consider which hardware to order.
The easiest way is to find a reputable dealer with an explicit warranty and replacement policy since this will solve many problems with the hardware. Most of the dealers also have a strict return policy that covers hardware for at least a year so we can count on getting the spare units we need.
In our experience networking hardware is probably the easiest to replace with the refurbished one. Networking equipment is something that usually simply works[6].
If we are working with the dealer that tests the equipment before delivery we can almost be sure that the equipment will function correctly.
Depending on the network speed and the required adapters, we will also order the cables end connectors but when it comes to the access network, the patch cables should probably be bought new since they are most prone to damage.
When it comes to exact configuration in our experience when buying servers the best strategy is to contact the dealer and find out what servers they have in stock. Once you know that, we can define the services that will run on the server, for example, storage, databases, or virtualization, and then we can decide on which particular server to buy.
Almost all the dealers offer customization of the servers they have and the main thing we will usually do is maximise the amount of memory in the servers.
The only thing that we are not going to buy refurbished is the disks since they are the most vulnerable part of any system being at the same time the part that gets the most wear and tear. Remember to always order trays and different components needed to fix the disks and other adapters to the servers. Some dealers will offer things like this for free, but always check with them before agreeing on the server configuration.
Other than servers, we can also buy refurbished UPS (Uninterruptible Power Supply) devices and power

equipment, but we should pay attention to the batteries since they are basically spare parts and should be bought new.

In the limited time we are running such an environment we still haven't managed to collect all the empirical data needed for a serious conclusion, but we can say a few things as preliminary observations.

We tried running a virtual environment that students use for testing network and software configurations. First iteration used 40+ servers which was later reduced to 8 servers in total. Storage servers are separate entities.

In regards of budget, preliminary results say that we created a data server that is about 40% of the price of the equivalent center that would be created from new hardware. Performance wise we are more or less at the same perceived speed. Only thing is that power usage was much (almost order of magnitude) greater when using the oldest servers. We still have no numbers to relate last generation of servers we are using to new equipment.

V. CONCLUSION

It is completely feasible to create a working data center using mostly refurbished equipment. Depending on the usage scenario this may have two big advantages. First is reusing equipment and being completely eco friendly by not throwing away perfectly functioning devices and second is the price. While some of the devices cannot be reused, such as disks, most of the datacenter can be refurbished. Market for refurbished servers exists and is quite large making it even possible to buy devices with warranties, although almost exclusively we are talking about being able to exchange devices by returning the malfunctioning unit and getting another refurbished one in return. If we utilize this model it is possible to save money and help the environment while attaining almost the same level of service for the users.

On the other hand, some usage cases cannot support this model of refurbished data center and almost require new equipment.

Considerations for purchasing or othervise acquiring refrubished equeptmen should be:

- Are we getting equipment that was checked before coming back to market?
- Is there a "replacement" warranty in place?
- Is the vendor able to provide multiple units (usually the signal that replacements will be available later)
- Is the equimpent tested?
- Can the vendor provide his list of references? Big refurbished equpment vendors can provide you with the actual source of the equipment.
- Is equipment delivered cleanly reset and upgraded to the latest versions of the firmware?
- Is vendor providing guidance and advice on purchases?

Although not exaustive, this list of questions should help when deciding on a particular vendor in a market that is increasing each year.

REFERENCES

[1] Bizo, D. (2022). Uptime Institute Global Data Center Survey 2022 Resiliency remains critical in a volatile world

[2] Gall, J., & Gall, J. (2002). The systems bible : the beginner's guide to systems large and small : being the third edition of Systemantics. General Systemantics Press.

[3] Bashroush, R., Rteil, N., Kenny, R., & Wynne, A. (2022). Optimizing Server Refresh Cycles: The Case for Circular Economy with an Aging Moore's Law. *IEEE Transactions on Sustainable Computing*, *7*(1), 189–200. https://doi.org/10.1109/TSUSC.2020.3035234

[4] M. Dayarathna, Y. Wen and R. Fan, "Data Center Energy Consumption Modeling: A Survey," in IEEE Communications Surveys & Tutorials, vol. 18, no. 1, pp. 732-794, Firstquarter 2016, doi: 10.1109/COMST.2015.2481183.

[5] Daigle, B. (n.d.). Data Centers Around the World: A Quick Look The Exponential Increase in Data. https://www.usitc.gov/publications/332/executive_briefings/ebot_data_centers_around_the_world.pdf Accessed on 28.4.2023.

[6] Sigcomm 11 Conference Committee. (2011). *Proceedings of Sigcomm 2011 and Best Papers of the Co Located Workshops.* Association for Computing Machinery.