



WHITEPAPER

Taking the Leap with On-Premises or Co-Location Datacenters

How and why companies are successfully building
and deploying their own infrastructure

In partnership with

DELLEMC

Introduction

In the beginning, there were server rooms: Office space dedicated to hardware housing valuable data and applications. These were expensive to build and manage, requiring dedicated IT staff and ongoing maintenance such as power and cooling.

For decades, these server rooms were the main solution for most companies. But then came the public cloud, and with it, cheaper options for storing and accessing data and applications.

In many ways, the public cloud was a revolution in how companies could develop and deploy products to customers. Containers, DevOps, Artificial Intelligence and Machine Learning, edge solutions — the arrival of the cloud brought with it a flood of these and other innovations.

But a funny thing happened on the way to public cloud dominance. Some companies found that they could be better served by going back in time to on-premises or co-location datacenters.

In this whitepaper, we're taking a look at this growing trend. Specifically, we will be examining and explaining:

- Why companies are going the on-premises or co-location route
- Considerations before building your own datacenter infrastructure
- What goes into on-premises or co-location datacenters

Let's dig in ...

Reasons for Going On-Premises

One driving force behind companies building their own datacenter is, somewhat ironically, cost.

Whereas startups and mid-sized companies can see major benefits to their budgets from cloud services, companies that reach a certain scale often find the math is no longer in their favor. At some point, the costs of storing and computing on the public cloud supercede the costs of running their own datacenter.

There is a tipping point many companies hit where it becomes more cost-effective to build their own on-premises solution. That tipping point scale slides up and down depending on the relative simplicity of an application or the use of high-value cloud services that require significant investment to build and operate. They will then either staff it themselves or work with a third party to manage it for them.

There are other reasons that drive interest in on-premises, including regulatory and compliance requirements. Cloud providers have quickly matured in this space and are now meeting stringent government security requirements.



Considerations Before Going On-Premises

The first consideration for companies exploring their own datacenter option is the investment, which can be substantial. One benefit, however, is that it will be a capital investment in most cases, which allows for some flexibility.

Another consideration is the time it takes to design, build, and deploy a datacenter infrastructure, along with potential disruptions to workflow during the process.

Perhaps most importantly, the company needs to take a thorough assessment of their current capabilities. Questions to ask include:

- Is your team at the point where they can be expected to build and operate datacenter infrastructure?
- If not, are you willing to invest in the IT staffing needed to handle not just the build but long-term operations?
- Have you undertaken a technical and financial assessment?

While the first two questions are important, getting an answer to the third one can be absolutely critical in order for a company to successfully deploy their own datacenter.



Solving technical problems requires tracking, measuring, and improving enterprise systems.

Companies that don't have an understanding of their technical maturity risk flying blind, which is why it's important to employ a thorough assessment such as our own [Technical Maturity Framework](#).

Through a technical maturity assessment, companies are able to:

- Align technology investments to business objectives
- Triage competing priorities
- Deploy emerging tech
- Measure progress with qualitative data

For major endeavors like building out an on-premises datacenter infrastructure, having the ability to do each of the above can be the difference between a wise and a wasted investment.

You can learn more about our Technical Maturity Framework by [downloading our free eBook](#).

What Goes into On-Premises or Co-Location Datacenters

In general, companies building out their own datacenter infrastructure begin by understanding the challenges they face. These can be things like scalability, latency, and security. They directly affect finding the right datacenter solution to fit your needs.

To be more specific:

- How much compute and storage capacity do you anticipate needing now and in the future?
- How much will latency be a factor in what you are trying to achieve?
- What are your security needs?

Once these challenges are addressed and understood, the design and build process can begin. This process usually involves four phases:

- 1. Solution Design**
- 2. Sourcing and Procuring**
- 3. Building Process**
- 4. Configuration, Testing, and Validation**

Let's break these down ...

1. Solution Design (approx. 1 month)

Designing datacenter infrastructure is not unlike tuning an instrument in some ways. Every component, every compute node and hardware selection, should be tailored toward specific needs.

Often times, an ODM provides hardware with unnecessary components that can be stripped out. A company's application may be memory intensive, which makes it possible for them to move away from the fastest processors. Workloads may be IO intensive, requiring an investment in the right disks for support. It's important to understand a workload requirement and tune each component to optimize performance and value.

Every decision made during the design process needs to be made with the company's specific challenges and goals in mind, from storage to compute to rack design.

2. Sourcing and Procuring (approx. 3-4 weeks)

Most datacenter solutions are multi-vendor due to the rarity of one single vendor being able to handle all the pillars that go into a datacenter.

Because one vendor may provide compute, another network, another storage, and so on, procurement can be a challenge to oversee.

Add in factors such as lead times and its effect on build schedules, along with delivery logistics and assembling, and the entire process of sourcing and procuring can take up to a month by itself.



3. Building Process (approx. 10 days)

Once design, sourcing, and procuring are finished, it's time for assembly. With all the materials landing at the location, they need to be unboxed and integrated.

If a company chooses to go with a co-location separate from their own premises, they will be paying for space, power, and cooling to the third party. And if they don't want to devote in-house engineering talent to the build, they will be investing in either Site Reliability Engineering (SRE) or an additional vendor for the build.

Each of these options can add considerable cost to the project, but again, if a company has reached that tipping point where on-premises is more cost-effective than continual investment in the public cloud, they should stay within budget.

As for actual on-premises builds, such as within a dedicated footprint within the company, additional costs to account for are power, cooling, and the actual space itself, as well as on-site personnel — or a third-party vendor — to oversee the build and ongoing management.

4. Configuration, Testing, and Validation

Included in the building process, configuration, testing, and validation is the minutiae of the on-premises datacenter solution. Many companies are not very mature when it comes to this phase of deployment, so the process is often done through a partner.

At Redapt, we've built our own system for configuration, testing, and validation named Forge. This automation tool, which has been in development for the past 10 years, connects to the network and creates a profile indicating all the settings and specifications a server needs. Automation will then take this profile and lay it over a hundred servers.

One benefit of going with a tool like Forge is that it greatly reduces the risk of human error during configuration. There's no need for provisioners to take an all-hands-on-deck approach on tight timelines, or hiring employees to have crash carts. Forge is designed to simply plug in and get to work.



Conclusion

The cloud, on-premises, hybrid — there's no one go-to solution when it comes to enterprise data technology.

For those companies looking to explore an on-premises datacenter solution due to cost or other circumstances, however, it's beneficial to go into the process with eyes open. That means:

- Understanding your technical maturity
- Clearly defining your goals and needs
- Researching the right solution for design, build, and configuration.

We hope this whitepaper has served as a good resource for understanding just what's involved in going on-premises with your data infrastructure.

Thanks for reading.



We can recommend an on-premises, cloud, or hybrid solution that fits your needs.

Reach out to one of our experts at redapt.com/contact