

The Cloud Multiplier: Strategies for Enterprise Stack Optimization

From quick capacity to robust disaster recovery,
here are 10 ways to extend and enhance your data center.

Table of contents

10 Ways Cloud Helps Optimize Your Application Stack 3

1. You want more capacity but not the capital costs 4

2. You need to support multiple AI use cases 5

3. Zero trust security is increasingly a necessity 6

4. Managing compliance with a lean staff is causing strain 7

5. The business says, “Flex” but finance says,
“Use what you have” 8

6. This sounds familiar: Our technology is special
and can’t possibly run anywhere else 9

7. You can’t remember the last time you fully—and successfully
—tested your DR plan 10

8. You struggle to adopt the latest technologies,
unlike some competitors 11

9. You have data, and it can’t travel—at least not easily,
securely, or inexpensively 12

10. Your private cloud is costing more than it was supposed to 13

10 Ways Cloud Optimizes Your Application Stack

Michael Chen
Senior Writer

Are you strategically supplementing your on-premises infrastructure with cloud solutions to achieve peak efficiency and agility? For many enterprises, the answer is yes. The Uptime Institute's [2024 Global Data Center Survey](#) reports that 55% of IT workloads are now run off premises and projects that by 2026 that number will reach 58%. In the Institute's 2020 survey, 42% of respondents' workloads were hosted outside their corporate data centers.¹

The goal for IT teams: Strike the optimal balance between cloud and on-premises resources. If you're not sure you have the proper fuel mix, ask: Are we over-buying hardware to provision for occasional peaks? Are we spending for special projects that will sunset before the equipment is fully depreciated? Do we have idle capacity because anticipated demand didn't materialize?

It's likely your experience has shown the cloud working well for quick development, failover, and applications with spiky demand. And it's often the place to start shopping if you need a new commercially available app or to replace outdated software. But cloud providers are improving their services to do and offer more, in more locations around the world. Calibrating the balance of systems running in the cloud versus on-premises needs to be an ongoing process.

Is it time to move more workloads? Here are ten signs the answer just might be yes.

1 You want more capacity but not higher capital costs

CapEx versus OpEx is perhaps the original argument for the cloud. Most organizations have set on-premises equipment upgrade cycles, but as new needs arise, getting additional capital allocations can be difficult and slow. You may not have planned to add GPUs to your data center. Networks might be overtaxed by new workloads that depend on AI and require storing far more data, so storage systems could need major upgrades. That all means new spending. But purchasing this hardware now might not be a wise use of capital.

The question to ask is: Will new investments deliver ROI beyond the immediate need? Will adding capacity to that storage array solve our problems for the foreseeable future, or would developing a data warehouse in the cloud let us do more? Are we ready to bring AI analytics tools and other technologies into our own data center—and ramp up the expertise to manage them—or will cloud-based AI services working in conjunction with a cloud-based data warehouse offer more sustained returns?

Capital spending in the data center may well make sense. But today, relatively small shifts in OpEx can satisfy short-term demand while proving concepts in a way that CFOs like to see. Budget might just follow, and in the meantime, you've learned what works for your needs.



2 You need to support multiple AI use cases

AI is a versatile technology that can be applied to specialized tasks: intelligent document processing, computer vision, chatbots using natural language processing, voice recognition and speech synthesis, predictive analytics, Internet of Things (IoT) automation, edge services for autonomous robots and drones, and increasingly, generative AI that creates text, images, and even applications. The use cases are there, and business units know it and are clamoring for support. Taking advantage, however, requires a wide variety of development tools, open source frameworks, libraries of pretrained models, and AutoML platforms—a diverse software component mix that's cumbersome to independently procure, configure, and maintain. And that's on top of expensive high performance CPUs, GPUs, storage, and networking infrastructure.

Hyperscale cloud providers are rapidly developing and delivering innovative AI technologies. They're also offering AI-assisted development environments that simplify using all your data in ways that truly inform the business. So it's not just about moving CapEx to OpEx. It's also about providing AI-driven innovations for the business faster and with less investment in hardware, software, and skills.

In the cloud, enterprises can flexibly provision compute resources, including large clusters of GPUs, as needs arise and start building and running AI without prohibitively expensive data center capital investments. Most cloud infrastructure providers put the tools and services needed to develop and deploy effective AI at their customers' fingertips.

The cloud can be a faster and less expensive path to creating systems that help make business processes more efficient, provide deep operational insights, forecast trends, judge market sentiment, and much more. And when it's time to build customer-facing applications, the cloud is often the place to do it. Deploying apps regionally for performance and to address local rules for data residency and other considerations also tends to be easier in the cloud.

3 Zero trust security is increasingly a necessity

Organizations are turning to [zero trust security](#) models across their networks as an effective means of keeping attackers away from their systems and data. If you've made that move, you know the benefits of granting the minimum access needed for users to complete their work. By abandoning the notion that perimeter protections are sufficient, you can help improve security while making it possible for team members and trusted partners to work with data wherever it makes sense.

Because zero trust needs to be end-to-end, some companies have hesitated to add cloud services, which may not have been compatible with the model. The good news is that more and more cloud providers do offer a zero trust security option, often with unique capabilities that not only limit access to data and systems but also hide details of the network, making it that much more difficult for attackers to compromise systems.

If the benefits of the cloud sound good but the need to extend your zero trust environment has kept you from moving applications and workloads, it's time to take a second look.

How Zero Trust Works

Assume no entity is trustworthy, and grant only necessary access with the least privilege possible. Segment the network to restrain any bad actor who may get in.



People



Services



Devices

Authentication:

Who are you? Use strong passwords and multifactor identification.



Context:

Where are you, and how secure is your device? Limit access based on contextual data.



Least privilege:

What's the minimum access needed to fulfill your request? Make time- and resource-limited access grants.

Microsegmentation of the network to the extent possible is a cornerstone of zero trust.



4 Managing compliance with a lean staff is causing strain

As data grows in volume and interconnectivity expands, complying with regulations keeps getting more complicated. IT teams and other groups, including legal, spend a great deal of their time and resources keeping up with GDPR, HIPAA, and other laws and regulations that affect IT infrastructures. That leaves fewer resources for the many other tasks involved in keeping a corporate data center—and a business—running at peak performance.

Hyperscalers offer region-specific data centers to address the common requirement to store data within a specific jurisdiction, with the ability to easily move data between regions and keep management and log data where it makes sense. Cloud providers also employ advanced tools, such as identity and access management controls, encryption, usage logging, and more, along with transparent documentation of up-to-date certifications and rule adherence, helping with audits.





5 The business says, “Flex” but finance says, “Use what you have”

Demand spikes happen, whether due to quarter end, a holiday sales boom, or an Instagram sensation. And when they do, you need to scale quickly or risk losing customers, sales, employee confidence, or some mix of the above. That’s where applications hosted on-premises face limitations. If you haven’t preemptively—and expensively—built in extra capacity to handle increased demand on monolithic applications, you may scramble to find money for hardware and IT staff overtime. This is especially true when responding to business demand means scaling up local applications, where any change may require a lot of institutional knowledge and testing time.

Cloud native applications that can be containerized, on the other hand, are designed to leverage the inherent scalability of hyperscale cloud systems. Because these modular applications let you add and remove resources quickly based on demand, you can realize a lower total cost of operations. When spikes subside, infrastructure scales back down, and you stop paying for that capacity until you need it again.

6 This sounds familiar: Our applications have unique needs and can't possibly run anywhere else

Almost every organization has monolithic applications running on-premises that require expensive, high performance, and finely tuned hardware configurations. In return, these systems manage the complex operations and processes intrinsic to the enterprise, such as finance, manufacturing, and supply chains. The awesome capabilities of these custom systems were built up over time, embody your organization's know-how, and deliver competitive advantages. Their uninterrupted operation is the foundation of your daily operations.

But when upgrades or integrations need to happen, changing the codebase is often a tricky business. You could update the application onsite, but you've already spent thousands of hours across years to customize and fine-tune it for your organization.

Bottom line, that application is stuck in your data center unless key parts of your infrastructure are also available from your preferred cloud. Top providers understand this and are responding. For example, you can now position Oracle Database services, including both Exadata and Autonomous Database systems, in any hyperscale cloud you like. That's one less limiter on your critical applications.

From monolithic to modular

Even if you plan to keep critical applications in your on-premises data center, it's well worth exploring a [cloud native](#) architecture. These platforms are designed to help applications deliver the same competitive advantages as monolithic software. They also offer an easier path for adding advanced functionality and are easier to secure and maintain. Moreover, modern cloud native applications can help reduce costs via automation, improve data analytics, and enable real-time operations.

They're inherently scalable and can run in public or private cloud infrastructures, or in hybrid and multicloud systems. Their use of containers, microservices, immutable infrastructure, and declarative programming tools allow engineers to make frequent, high-impact changes with less effort and ease integration with generative AI and other emerging technologies.

7 You can't remember the last time you fully—and successfully—tested your DR plan

Backups are easy. Actually recovering systems and data? A lot harder. Compared with hyperscalers, owned data centers are more constrained by physical limitations, including access to redundant systems, power and connectivity outages, and exposure to natural disasters. From a business perspective, regularly and fully testing BCDR plans can be disruptive and time-consuming.

Maintaining secondary facilities is expensive. As a result, many companies spend a lot of time creating detailed plans but do only piecemeal or tabletop tests. That's a risky business.

The cloud can help minimize downtime and data loss while, importantly, balancing the cost to protect each workload with its criticality. Hyperscalers do this by offering a variety of disaster recovery solutions, from simple data storage to standby systems to all-inclusive production environments. Cloud-based services also often come with their own one-click backup solutions that offload the administration and maintenance burden from your staff while providing a vetted solution.

Companies with hybrid setups can lower disaster recovery costs by creating so-called “pilot light” IT deployments—that is, caches of live, up-to-date data and standby services that can be used to restart a system in a cloud data center. It's major peace of mind that is very expensive to accomplish with only traditional applications in owned facilities.

8 You struggle to adopt the latest technologies, unlike some competitors

Because of the scale at which they buy, hyperscale providers generally get first dibs on new technologies—including the scarce and pricey GPUs required for AI—and they deploy them quickly, giving their customers an advantage over rivals that must buy, configure, and test new technologies in their own data centers.

If your competitors are continually seizing an advantage by running the latest and greatest innovations, they're likely either dedicating a lot of time and money to do so or accessing that tech via cloud services. Working with a cloud provider can help lower the time to value and takes the work of getting the infrastructure just right off your plate.





9 You have data, and it can't travel—at least not easily, securely, or inexpensively

For many enterprises, even those making use of the cloud, the data center is the center of data gravity. When it comes to using data for AI, analytics, or other initiatives, that may limit their options, whether for reasons of latency, security, or practicality. Partnerships between cloud providers can allow for accessing data in one cloud from another by using a proprietary interconnect to reduce latency and egress charges. But that's complex and not always sufficient.

Fortunately, a new type of partnership is changing that calculation for Oracle Database customers looking to run workloads in the cloud without sacrificing format, customizations, or security. They can now use their Oracle databases, running on Exadata systems if they choose, to inform applications in the hyperscale cloud or clouds of their choosing—AWS, Google Cloud, Microsoft Azure, or Oracle Cloud Infrastructure (OCI)—and assemble a diverse yet unified cloud portfolio without incurring data egress fees or the performance hit imposed by interconnects.

Organizations also gain simplified budgeting and billing and improved data governance and security while achieving the low latency previously possible only by connecting databases to applications in the same physical data center.

10 Your private cloud is costing more than it was supposed to

Many enterprises are all in on VMware—Broadcom reports that 85 million workloads now run on the platform, and 5 million developers actively build on it. Millions of applications and databases have been optimized to run on VMware, and corporate IT teams have mastered every feature and function. It wasn't so long ago, in fact, that VMware implementations were thought of as the alternative to public cloud.

That has changed in recent years. Hyperscale cloud infrastructure providers, one by one, have partnered with and embraced VMware. But not all offerings are created equal.

It's important to compare capabilities in the cloud because providers may not fully support VMware exactly as it runs in your own data center. Those that added VMware as an afterthought often limit your access to administrative functions. Accessing the capabilities of these clouds may mean you need to change your processes, and third-party software that you have licensed for your VMware environment might no longer work.

In contrast, services such as [Oracle Cloud VMware Solution](#) allow vSphere compute virtualization, NSX networking, and vSAN storage to be rapidly provisioned and scaled, consistently secured, and easily managed on a pay-as-you-go basis. Oracle's joint offering with VMware goes even further by allowing granular control of these environments for version management and patching and the ability to manage VMware infrastructures using native VMware tools. These on-demand deployments connect to and extend on-premises VMware environments, making it easier to incorporate cloud-based capabilities, such as AI and advanced analytics, into critical business processes.

Oracle helps you expand beyond your data center

[Oracle Cloud Infrastructure \(OCI\)](#) can help your business, whether you want to stay on-premises, move workloads to the cloud, or use multiple clouds. Oracle's distributed cloud gives you access to more than 150 cloud services in your location or in more than 50 global regions, and [Oracle's multicloud](#) solution offers Oracle Database on optimized hardware in all hyperscale clouds.

Oracle offers Oracle EU Sovereign Cloud to help address EU data privacy and sovereignty requirements; government cloud regions in the United States, UK, and Australia for operational and defense customers; and [Dedicated Regions](#) that put the full suite of OCI services in your preferred location.

Confidently move your existing workloads as is to OCI, which enables an on-premises experience in the cloud, including optimized hardware for Oracle Database. Take advantage of cloud native functionality at your convenience, both for existing applications and new use cases, with consistent pricing across all regions and deployment options. Leverage lower-cost data egress to move data where you need it, including to other clouds. Gain early access to advanced AI technologies that, as of January 2025, include the largest AI supercomputer of all the hyperscalers.

OCI enables you to augment, expand, or exit your data center. It's your choice.

Learn more about OCI

Try OCI for free

Connect with us

Contact us at oracle.com/apac/contact/

Copyright © 2025 Oracle, Java, MySQL and NetSuite are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners. This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

