WHITE PAPER

DIAMANTI

Build vs. Buy: Choosing the Right Kubernetes® Infrastructure

Kubernetes adoption is accelerating rapidly. A recent survey by the Cloud Native Computing Foundation (CNCF) showed that 78% of over 1,100 respondents are using Kubernetes in production, a huge jump from 58% last year[1]. Existing IT infrastructure has been optimized over a period of many years for virtualized business applications and may not efficiently support containers and cloud-native applications. As your organization navigates the transition from virtualization, you'll need an infrastructure that addresses the unique needs of Kubernetes.

Given the increasing pressure on IT teams, you'll also want to carefully assess whether you will buy or build your infrastructure to support Kubernetes. As with any infrastructure decision, there are a number of factors to consider:

- **Solution components:** What components—hardware and software—make up the solution, and how well do they meet your requirements?
- **Deployment time:** How long will the solution take to deploy, and how much expertise will that require? Solutions that exceed your team's skills may need expensive professional services engagements, adding cost and time.
- **Management:** How much time and expertise will it take to manage the various hardware and software components?
- **Scalability:** How difficult is it to scale the solution as your needs grow?
- **Total cost of ownership (TCO):** How much will it cost to own and operate the solution, including staff time?
- **Support:** Who would you contact to get technical support when problems arise?

9 out of 10 survey respondents who are using containers are doing so in production.

"By 2023, more than 70% of global organizations will be running more than two containerized applications in production, up from less than 20% in 2019[2]."
–Gartner

Diamanti is Kubernetes certified

[1] Cloud Native Computing Foundation (CNCF) "Deployments are getting larger as cloud native adoption becomes mainstream," March 2020

[2] Gartner "3 Critical Mistakes That I&O Leaders Must Avoid With Containers," Jeffrey Hewitt, August 2019

For Kubernetes environments, there are a few additional factors to consider:

- **Bare-metal or virtualized?** Kubernetes was originally based on Google's Borg project which was designed for bare metal. Running containers inside virtual machines (VMs) adds an additional layer to the stack that must be managed and debugged, adding cost and complexity.
- **Persistent storage.** Applications running inside containers need to be able to save data permanently. Container solutions must provide a mechanism for persistent storage, even as containers come and go.
- **Networking model.** Getting networking right remains one of the most difficult aspects of container environments, and container networking must integrate smoothly with your existing data center networking.
- **Support.** How will you get support for the full infrastructure stack, including both hardware and software? From a single vendor? From multiple vendors? From the open-source community?

This paper examines the tradeoffs between available solutions for running Kubernetes on-premises.

	Do It Yourself (DIY)	Virtualized Infrastructure	Diamanti
Bare metal?	√		√
Integrated Kubernetes		√	√
Integrated compute, storage and networking		√	√
Full-stack 24×7 support		√	√
Hardware and software updates and patches		√	√
Real-time I/O services			√
I/O performance	Variable	Low	High
Infrastructure efficiency	Medium	Low	High
Time to deploy	Months	Weeks	Minutes
Management complexity	High	Medium	Low
TCO	High	Medium	Low

Do It Yourself

The biggest advantage of a DIY Kubernetes solution may also be its biggest disadvantage: complete freedom to choose all the hardware and software components.

- Do you want servers with only internal storage, or servers with separate storage arrays?
- Bare-metal or virtualized?
- What version of Linux?
- Docker or a different container environment?
- What networking plugin?
- Open-source or packaged software distributions?
- What other software (drivers, etc.) is needed for everything to work together?
- How will you monitor and manage the environment? Through CLI only or do you need a UI?
- How will you secure the environment?
- How will the resulting solution integrate with existing data center infrastructure?
- How will you get support?

These questions aren't impossible to answer, but if your team is new to Kubernetes—or just needs to get a project off the ground quickly—there can be a lot to wrestle. In today's rapidly evolving IT environment, the question of buy versus build has become a critical one.

If you choose to build out infrastructure yourself, the initial deployment will take longer and may require professional services to complete. Ongoing management complexity will also be higher than it would be with a more integrated solution since you'll have to keep up with the open-source community and integrate patches and enhancements as needed.

If you opt for servers with internal storage, scaling requires adding additional servers. However, you will require lifecycle management tools to seamlessly add capacity without impacting existing workloads. You'll also have to provide a mechanism for data management and data protection on each server, figure out how to balance storage use across the set of servers, and you may need to provide a mechanism for shared storage such as NFS or clustered file system such as Ceph or GlusterFS.

If you choose an external storage array, it can simplify storage management initially, but scaling becomes more complicated. Kubernetes can support hundreds of worker nodes which means you add servers until the storage array runs out of performance or capacity. Then you either add a second storage system or replace the original system with a more powerful one. Either way, this can be disruptive and can result in a big, and often unanticipated, incremental expense.

PROS

- Complete flexibility

CONS

- Longer time to deploy
- Management overhead
- Lack of full-stack support
- Must keep up with community
- Professional services may be needed
- Solution complexity
- Solution resiliency
- High total costs due to management/staff overhead

THE RISK OF "INTELLECTUAL" LOCK-IN

IT teams that take a do-it-yourself approach to Kubernetes may not always recognize the added Day 2 risks they face.

A talented team can create a tailored solution that meets all of the business requirements, but the ongoing maintenance and improvements to that system will require prior institutional knowledge and excellent documentation. If the original architects of the solution move on to new positions or new companies, the business may struggle making configuration changes or enhancements.

Do-it-yourself infrastructure is like getting a puppy. The real work begins with care and feeding.

When it comes to networking, you're completely on your own. You need to make sure that your chosen networking model is compatible with your container and orchestration solutions. Again, that may not sound difficult, but architecting the solution and selecting the right solution can add time and complexity.

Finally, there is the issue with technical support. With a DIY approach, troubleshooting issues may require contacting several different vendors. If you use a commercial Kubernetes distribution, you may be able to get support for any software issues, but if the issue is the underlying hardware, you'll need the different vendors to collaborate on the issue.

Converged and Hyperconverged Infrastructure Solutions

A variety of vendors have created converged and hyperconverged infrastructure solutions to reduce the complexity of IT infrastructure deployment:

- Converged infrastructure (CI) pre-packages several servers with a separate storage array.
- Hyperconverged infrastructure (HCI) combines servers with internal storage, software to virtualize that storage.

CI/HCI solutions have traditionally included virtualization software, however there is also a new category that provides a similar integrated approach with bare metal.

STANDARD VIRTUALIZATION-BASED CI/HCI SOLUTIONS

These solutions simplify hardware deployment, but require you to run Kubernetes on top of the hypervisor. The additional layer of abstraction adds unnecessary overhead and impacts overall performance. In a virtualized environment, each VM has a full blown operating system. Deploying Kubernetes on bare metal means applications are isolated at the operating system level, allowing applications to access hardware without the need for pass-through or hardware emulation. By sharing a single operating system, organizations can see higher consolidation levels, have fewer operating systems to patch and maintain, and save money from OS and virtualization licensing.

After initial deployment, scaling is relatively straightforward for HCI since you simply deploy additional VMs configured with the full software stack on new servers as they are added. With CI, you face the challenge of scaling storage separately, a larger incremental cost.

Container networking in virtual environments relies on two levels of routing as traffic is routed to the right virtual machine first, and then to the right pod on that virtual machine. The worst-case scenario results in deployment of a single container per VM, wasting resources and eliminating the density advantage of containers.

PROS

- Multiple consumption options

CONS

- Virtualization licensing costs and higher management overheard drive up TCO
- Performance is reduced
- Bewildering range of choices
- Complex software deployment
- Virtualization plus network limitations reduce container density

With virtual environments, each OS instance increases the memory and storage footprint, thus affecting the available storage not only for your applications but also for storage services such as backup and replication. This increases the administrative time and costs of managing storage.

Traditional CI and HCI solutions are designed for virtualization rather than containers, making bare-metal container deployment impossible in almost all cases. The vendors themselves remain largely focused on virtualization which means you have to rely on the open-source community for container and Kubernetes support.

Diamanti Hyperconverged Container Platform

Diamanti has created the first infrastructure platform purpose-built for Kubernetes, combining the power of hyperconverged infrastructure with the unparalleled performance and efficiency of a platform that has been re-engineered specifically for distributed modern applications.

The Diamanti platform is the first and only Kubernetes solution integrated with a patented I/O-optimized architecture, delivering transformational application performance. With Diamanti, Kubernetes becomes an out-of-the-box solution, allowing organizations to focus on deploying modern applications across on-premises and hybrid cloud infrastructure and driving immediate business outcomes.

Efficient, shared infrastructure for containers provides persistent storage and plug-and-play networking that integrates with existing data center technologies. The result is a highly-available pool of CPU, memory, network, and storage resources delivered to containers on-demand. Scaling occurs through the addition of nodes to a cluster.

Additionally, the solution integrates I/O offload cards that shift networking and storage traffic management off of the CPU. This unique architecture is designed for the demands of modern distributed applications that have higher levels of East-West traffic. By managing this traffic, the Diamanti platform can also deliver full quality-of-service (QoS) for all resources including storage and networking, guaranteeing application performance without code changes or customization. The Diamanti architecture delivers order-of-magnitude latency improvements compared with traditional shared-storage systems and software overlays.

The Diamanti platform integrates everything—hardware and software— out of the box, so it can be fully deployed and operational in minutes. You'll be able to start running Kubernetes immediately, without having to spend weeks or months standing up a DIY solution. Open-source software, including Docker, CentOS, and Kubernetes, is pre-installed and ready to run modern applications, so there's no vendor lock-in. An intuitive UI makes managing and monitoring the platform simple, even for those without prior infrastructure management experience.

PROS

- Full-stack integration
- Full-stack, 24×7 support
- Fast to deploy
- Easy to manage
- Easy to scale
- Eliminates vendor lock-in
- Highest container density
- High resource utilization
- Ultra-low latency

CONS

- May have higher capex than some DIY options

Only Diamanti offers:

- Guaranteed performance across memory, compute, storage and networking
- 10x higher performance (1 million+ IOPS) vs. traditional solutions
- Low latency and fast throughput
- No time spent configuring infrastructure for Kubernetes
- 6x improvement in infrastructure utilization vs. traditional solutions

Making Smarter Infrastructure Choices

The need to provide new digital services is forcing enterprises to pivot to containers and Kubernetes as the common platform for modern applications. As you make this transition, you have to consider your requirements and make careful infrastructure choices to avoid being saddled with a solution that is overly complex, difficult to manage, that lacks the necessary performance, or that locks you into a specific vendor environment.

Only Diamanti satisfies these needs. Full-stack integration and full-stack support allows developers and operators to avoid vendor lock-in while taking full advantage of familiar container software including Docker and Kubernetes. Storage and networking QoS deliver the performance necessary for demanding production environments. Diamanti requires no code changes. Using services that are already supported by standard Linux removes the need for custom protocols and drivers.

ABOUT DIAMANTI

Diamanti delivers the industry's only purpose-built, fully integrated enterprise Kubernetes platform, spanning on-premises and public cloud environments. It gives infrastructure architects, IT operations, and application owners the performance, simplicity, security, and enterprise features they need to get cloud-native applications to market fast. Diamanti provides the lowest total cost of ownership to enterprise customers for their most demanding applications.

Based in San Jose, California, Diamanti is backed by venture investors ClearSky, CRV, Engineering Capital, Goldman Sachs, GSR Ventures, Northgate Capital, Threshold Ventures (formerly DFJ Venture), and Translink Capital.

For more information visit www.diamanti.com or follow @DiamantiCom.

THE END OF MOORE'S LAW

For decades, the technology industry has relied on the ever increasing computing power delivered by CPU improvements—the doubling of processor power every two years, known as Moore's Law. However, since 2015, this rate of improvement has slowed down significantly as we have reached physical limitations. This, in turn, means that CPU power is becoming a premium.

At the same time, anywhere from 10 to 50 percent of a CPU's available processing power is handling I/O traffic in a server. This computing power is not available for running the actual applications, further limiting the capacity of a single server.