OPEN CLOUDS

POWER INNOVATION, AGILITY AND EFFICIENCY

SUPERMICRO

intel

EXECUTIVE SUMMARY

Cloud computing has come a long way in the past decade. Organizations that are moving their workloads to a cloud computing environment have several choices to make. Perhaps the most important is what kind of cloud best fits their needs. One of these considerations is whether to use a closed and proprietary cloud or a cloud that is open and based on industry standards.

Open clouds allow for a wide range of both hardware and software choices compared to clouds that are created and maintained by external organizations. An open cloud computing environment is possible today and can be constructed to be efficient for various workloads. Also, by implementing an open based cloud environment, applications can return results quickly. Applications that require fast responses and low latencies can take advantage of the latest innovations that can be easily incorporated into an open cloud computing environment. By creating and using a cloud that is based on widely available standards still allows for customization and can reduce costs over time.

This whitepaper will examine several factors that should be considered when investigating what an open based cloud computing environment looks like and how costs can be reduced.

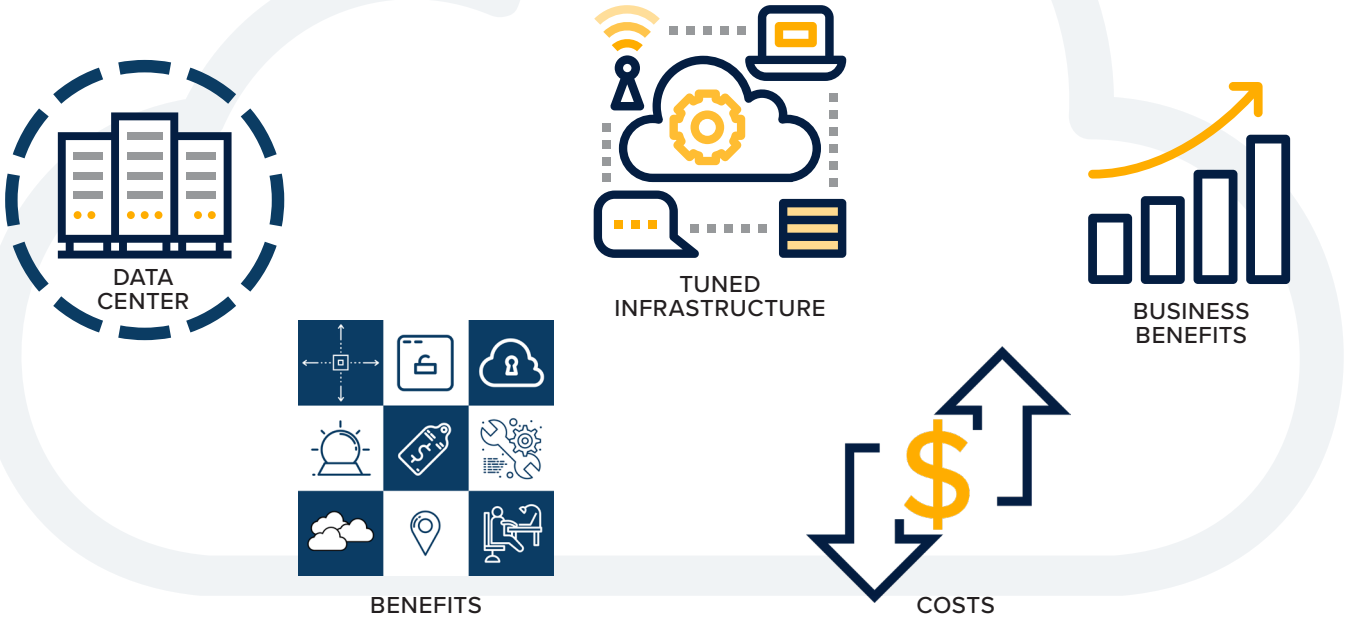

INTRODUCTION

There is overwhelming momentum and popularity around the move to the Cloud. Cloud architectures provide tremendous benefits in cost, efficiency, scale, time to market, and innovation. The next question is, "What type of cloud do I need?" The critical distinction here is Open versus Proprietary Cloud, and not the false choice of Public Cloud versus Private Cloud.

Cloud computing has become an important topic and area of significant spending in the mind of IT organizations. The relatively inexpensive (historically) availability of computing resources, storage resources, and services has led most IT organizations to move some or all of their workloads to cloud environments. While there are many choices to be made when an organization decides to embrace cloud computing, one of the most important is what kind of cloud environment is the most efficient in terms of servicing their users as well as reducing costs.

The National Institutes of Standards and Technology (NIST) defines cloud computing to be (Reference 1):

Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. This cloud model is composed of five essential characteristics, three service models, and four deployment models.

The NIST defines various deployment models for cloud computing, including public, private, community, and hybrid clouds. However, one of the most important considerations is whether the multiple components of a cloud computing environment are proprietary or are built and implemented on open systems and software technology.

An on-premises cloud is a type of cloud computing where the physical infrastructure that includes servers, storage, and networking, reside within the company's real estate. There are benefits to building and running this type of infrastructure to service the users. Still, the main question that needs to be considered is whether an on-premises, physically-based cloud allows for the use of the latest technology and can be scaled according to workloads.

OPEN STANDARDS DATA CENTER

A data center is designed and constructed to serve a wide range of users needs to be responsive when required, scalable as workloads grow, and to take advantage of new technologies when available and appropriate.

> *"There are two ways to look at how 'open' is defined," explains Mike Evans, VP, Technical Business Development at RedHat. "One is to look at how open the software is that's used to build the cloud, while the other is examining the software that's put on top of the cloud."*

Evans explains that most software used to build the cloud, whether Kubernetes or another kind of software, is open source. Only Microsoft builds its cloud on a proprietary engine (Windows), and even then, it's being moved over to the open-source Linux software.

The concept of using open source software by a proprietary cloud provider has become an issue for organizations that genuinely contribute to the community. Although some cloud providers contribute to the open-source community, there is evidence that cloud providers have taken open-source software and, with no contribution back to the community, offer a service and charge money for this service (Reference 2). This has caused several vendors to modify licensing terms so that an open-source software product cannot be monetized so easily.

Most software that sits on top of the cloud, such as containers or developer tools, are open-source operations as well, with one important caveat: there is no "open source of everything option," as Red Hat's Evans puts it. "When you move to private and hybrid cloud systems, you have more options, but you're still going to encounter a legacy, proprietary element at some point. This proprietary layer is probably going to happen at a level where no organization would need or want to be making changes. Most platforms today are examining interoperability options with other platforms, which is where the desirable "openness" is essential."

The components that go into an efficient and high performing cloud include:

HARDWARE — While there are several choices in terms of the CPUs for compute servers, the dominant CPU in use today uses the x64 instruction set. A leader in this category are 2nd Generation Intel® Xeon® Scalable Processors, Powered by Intel®. This standardization allows for a wide range of applications to run without modification. However, since there are several choices of CPUs that fit this high-level requirement for a specific instruction set, and the different options can make a sizeable native performance difference. Although virtualization and containerization technologies can abstract the underlying differences, the optimal matching of a CPU to the application will result in increases in performances and potential decreases in energy consumption.

SOFTWARE — The software stack required for a smooth-running cloud environment can be complicated and highly specific for individual mixes of workloads. The underlying libraries and management software requirements are almost guaranteed to be different from company to company. Without a wide range of choices that can easily be installed and configured on the underlying hardware, a cloud may not serve the needs of the users or system administrators. Not all middleware and supporting software will run optimally on all CPUs. Choices abound for all layers in the software stack, and an open computing environment is key to creating an efficient cloud computing system.

NETWORKING — Many of today's most innovative applications require clusters of servers, sometimes working in coordination with each other to solve a complex problem. In other scenarios, different servers are used to perform more simple tasks, with each server being given a certain amount of work to do, completely independent from other servers. The networking between servers needs to be matched to the application requirements, both in latencies and bandwidth. An open based cloud service needs to be designed and implemented with the networking requirements defined. Depending on the workloads and applications being used, different networking solutions may need to be used, and the application should not be locked into an environment that is not optimal.

COMPONENT	OPEN	PROPRIETARY
SERVERS	Many choices	Only what vendor decides
STORAGE	Many choices	Only what vendor decides
NETWORKING	Many choices	Only what vendor gives you access to
SOFTWARE STACK	Create what is needed	Only what vendor certifies

TUNED INFRASTRUCTURE

Modern and responsive applications require a tuned infrastructure. There are many tuning parameters that, when understood, can lead to a more productive and lower-cost cloud computing environment. Servers can be “tuned” with changes to the BIOS based on application signatures. While the absolute performance of a given server based on an offered CPU clock rate may not be exceeded, most applications do not achieve the raw performance rate. System manufacturers provide some tunable parameters to the system administrators, and these parameters can be tuned for the application that is to be executed. Specific parameters can lower energy usage, increasing efficiency, and reducing costs.

Networking is another infrastructure component that can be tuned in an open cloud computing system. As mentioned previously, networking can have an enormous impact on the performance that the end consumer will notice. With different networking technologies working in tandem or individually, selecting the appropriate options for a known workload can significantly affect the performance of the entire system.

Storage is overlooked in many cases when deciding on a cloud architecture. While the focus for some architects is the servers and CPUs, without choosing the correct and optimal storage service, there can be an undesired effect on all applications. There is a wide range of storage options available today that can satisfy many application requirements. An open computing cloud environment that is servicing the needs of many applications running concurrently may need to use a range of storage options — merely assuming that one size fits all for tens to thousands of applications will ultimately lead to these applications not serving their users effectively.

BENEFITS

There are significant benefits when implementing and using an open based cloud environment. More and varied applications can run with a broader range of choices for the underlying hardware and software. Open-source applications can be used and installed easily without paying proprietary providers a fee for using open-source software.

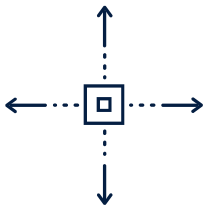

Flexibility

An open based cloud gives the implementers a wide range of choices regarding the hardware and software. For example, many choices are available for servers, ranging from one to eight sockets and up to many terabytes of internal memory. The available options once the number of sockets in a server is determined, the combinations of the number of computing cores, high-speed memory, internal disk capacity and type, and number of accelerators installed is extensive. The packaging can also be chosen, whether rack-mounted, blades, or more optimized enclosures. Open based systems allow for combinations of the components so that a proprietary cloud would not allow. Specific components can be selected and used depending on the workloads. A hybrid cloud approach, if a public cloud is required for some workloads, allows for maximum flexibility of the portion that the data center operator owns and controls.

Open Standards

Open standards are typically developed by several cooperating organizations. In the software area, the goal is to build software that can be implemented on various hardware systems. The best known of these efforts is open-source software, where developers can submit enhancements, bug fixes, and complimentary software. The software that is developed can be downloaded for free, or binaries can be widely distributed. By using open standards within a cloud environment, more applications are available to users. Within the hardware realm of open standards, the Open Compute Project (OCP) (Reference 3) has evolved into a global organization reimagining computer hardware design, reducing costs, and creating scalable solutions.

> *"I describe openness as the ability to choose the best hardware and software from a variety of options," explains Open Compute Project Foundation CTO Bill Carter. "Specifically, I mean hardware that is interoperable and available from many suppliers, and supports industry-wide interfaces and API's." Datacenter IT executives can now craft a data center that meets their needs more closely.*

Security

The security of data stored in a cloud environment has become one of the most critical concerns today. Many organizations limit where company or customer data can be stored, either within the "4 walls" of the enterprise and/or within specific geographies. The challenge of protecting sensitive data is two-fold. The first is the software security and hardware security that limits attacks, intrusion, passwords being compromised, or information being stolen. The second is the physical location that the data is stored. Countries, states, and economic zones may each have their laws and policies regarding the location of its citizens' data. Only an open cloud that takes advantage of the required technologies can satisfy many of the requirements. Relying on a proprietary cloud with limited physical locations and optional and unknown security practices could lead to data breaches. An enterprises data can easily reside in-house within the security setup (cyber and physical) that an organization needs.

Future Proofing

New compute, storage and networking technology continues to be developed at an astounding pace. An open cloud based on known and open standards can easily take advantage of these new technologies. For example, when a new GPU is available, the software that controls the GPU will already have been made available for systems based on open standards. An open cloud will incorporate this new technology earlier than having to wait for engineers to make the necessary changes for a closed cloud. This concept applies to more than just a new compute capability, but includes new types of storage, storage software, and networking. With new technologies constantly being introduced, an in-house, open cloud can quickly integrate these new capabilities as applications demand.

Transparent Pricing

When operating a cloud based on open standards, the owner/operator can easily decide what to charge clients for and what not to. Compared to a proprietary cloud-based where the operator wants to locate the physical facility, charges for data movement (ingress and egress) can be eliminated or significantly reduced. Since, in many cases, an open cloud will be located on the property of an enterprise, these data movement charges do not make sense and can be eliminated, reducing the overall costs to user organizations. Internal, open clouds can charge various departments for resource usage, without the complicated process of paying an external vendor.

Increased Serviceability

Clouds based on open standards are well understood and can be kept up and running more manageably than a proprietary cloud. Many companies can service their servers based on industry-standard components since the CPUs, memory, and peripherals have much in common from one vendor to another. Administrators can determine where an issue is taking place and the nature of the potential problem. To perform these actions, administrators rely on several tools, most of which are open-sourced and easily installed. By using these tools, cloud data centers can respond to users' needs more quickly, reducing downtime and increasing productivity. Internal or contracted service experts can more easily fix issues by using open source based tools, compared to relying on the operators of a public cloud.

Multi and Hybrid Cloud Options

Implementing a cloud can be either within an enterprise's facilities or at a commercial cloud provider facility. Modern organizations will most likely rely on a combination of cloud implementations and delivery mechanisms. Simultaneously, workloads will vary over time, whether daily, weekly, or seasonally, so it does not make sense to create a cloud data center for the maximum workloads. Instead, external and proprietary clouds can be used for bursting when an open cloud does not have the capacity to handle these bursts. A hybrid-cloud or multi-cloud approach that combines daily use, which is contained within an on-site open cloud, and occasional use of an external cloud will allow companies to optimize costs and remain responsive. Multi and hybrid clouds increase the flexibility of IT managers to match application use with required resources.

Location-Based Applications

A location of a cloud data center can have a substantial effect on the responsiveness of applications. Applications that need very low latency responses will perform better when housed in an on-site data center, as network latencies will be reduced. Fast database access or interactive applications where the user notices a higher latency in obtaining or accessing data can have a devastating effect on overall satisfaction and the ability to get work done. An internal cloud within a corporate data center gives users more predictable responses with lower latencies from running applications. Internal networks are more reliable and faster than requiring data to be transmitted over the public internet.

Optimized Workload Management

An additional consideration when designing a cloud is how workload management will be implemented. There will usually be competing departments requiring access to valuable compute and storage resources. An efficient and tuned workload management system that reflects priorities within an enterprise will reduce resource competition and deliver results faster. Utilizing valuable resources ultimately leads to a reduction in the Total Cost of Ownership (TCO). Workload management systems based on easily definable and open standards will lead to a more optimized cloud environment, reducing costs and increasing productivity. Different organizations within an enterprise may have different deadlines or priorities. By using an in-house open based cloud, these different workloads can easily be assigned different priorities, increasing the efficiency of a data center.

REDUCED COSTS

Costs to the consuming organization can vary depending on many factors. Overall, prices will be lower when creating or using a cloud based on open standards than a cloud based on proprietary technologies. This is due to:

Wide range of infrastructure choices:

Servers, storage, and networking based on open standards are widely available from several vendors. With competition comes decreased costs to the end-users. Also, software costs are reduced since software vendors can offer solutions that are already available and do not need customization for a proprietary cloud implementation.

Only buying what is needed:

When implementing an open cloud, whether on-site or elsewhere, the operator can purchase the required infrastructure based on workload requirements. External cloud providers typically over-provision, indirectly passing on costs of unused systems to the customer.

Movement of Services:

Open clouds allow for easy movement of applications and other services. When the applications are based on open computing platforms, it is relatively straightforward to move to another, compatible cloud. Standard-based infrastructures are designed to run standard-based applications. Movement of a critical computing environment may be required for many reasons, and establishing a cloud environment makes sense.

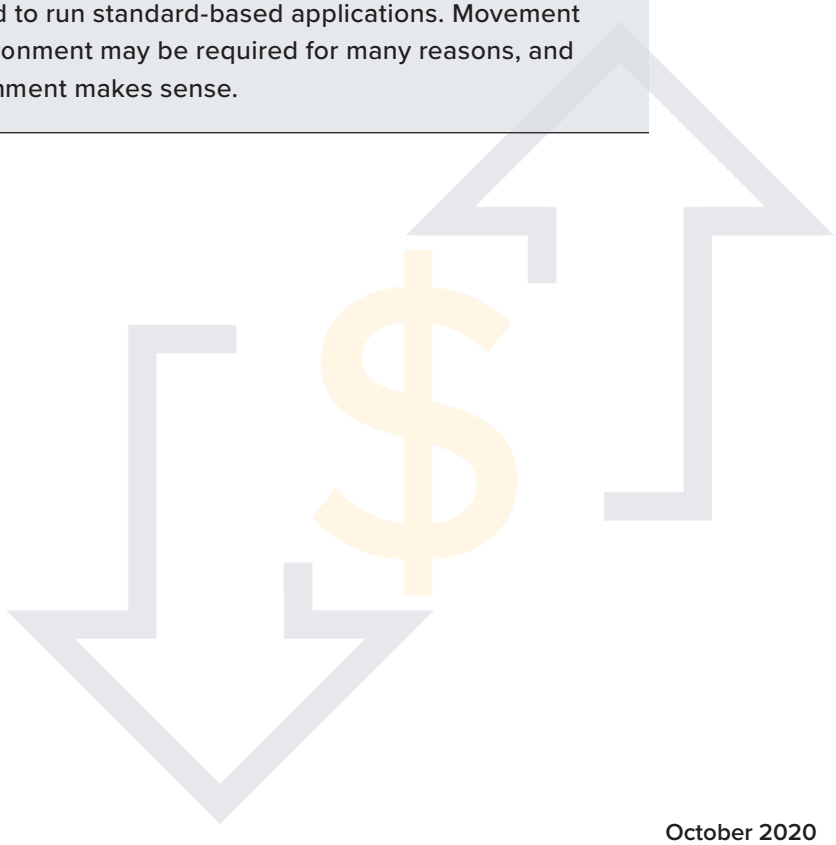

OPEN CLOUD BUSINESS RESULTS

An open cloud environment has numerous benefits that result in:

BETTER EFFICIENCY — Infrastructures become more efficient in terms of energy used and tuned more closely to the workloads.

MORE FLEXIBILITY — An open cloud can be adapted quickly to a changing workload.

CUSTOM TUNING — Hardware and software can be tuned so that the applications will run faster, consuming less energy

MODERNIZE AS NEEDED — Operators of open cloud data centers can upgrade to the latest infrastructure as needed, without relying on external schedules

COSTS AT KNOWN — Organizations can budget based on known costs and not be hit with unexpected charges.

INCREASED RETURN ON INVESTMENT — An open based cloud increases the ROI by only acquiring the infrastructure needed and keeping systems busy.

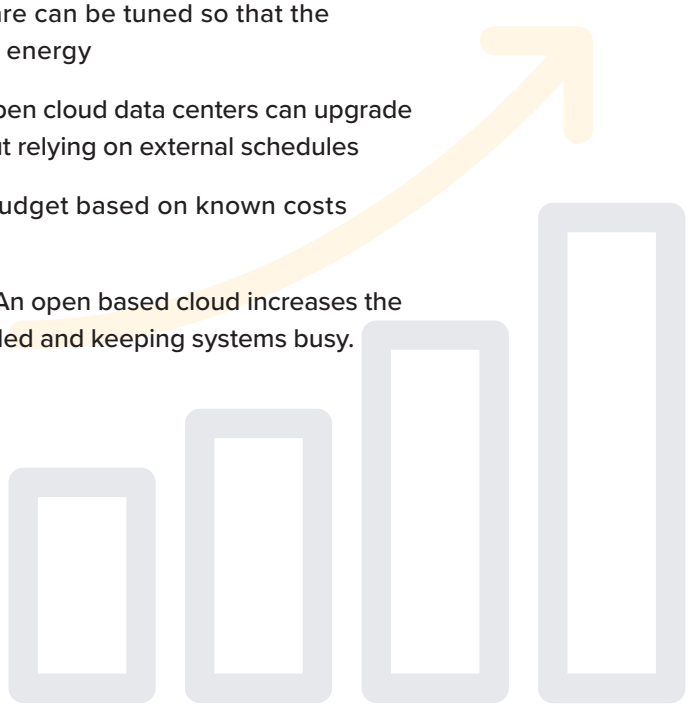

SUMMARY

Clouds come in many forms and delivery mechanisms. While there is much discussion about which form of a cloud to use, the higher-level concern should be whether to implement a cloud-based on open standards or use a proprietary one. The cloud computing market, estimated to be $ 364B in 2022 (Reference 4), is evolving. Different enterprises will need to determine what is important to them, not just today but moving forward as well.

REFERENCES:

1 - https://csrc.nist.gov/publications/detail/sp/800-145/final

2 - https://www.zdnet.com/article/amazon-and-commercial-open-source-in-the-cloud-its-complicated/

3 - https://www.opencompute.org/

4 - https://www.gartner.com/en/newsroom/press-releases/2020-07-23-gartner-forecasts-worldwide-public-cloud-revenue-to-grow-6point3-percent-in-2020

Super Micro Computers, Inc
980 Rock Avenue
San Jose, CA 95131 USA

Learn more at www.supermicro.com