WHITE PAPER

Protecting the Next Generation Datacenter: Hitting a Moving Target

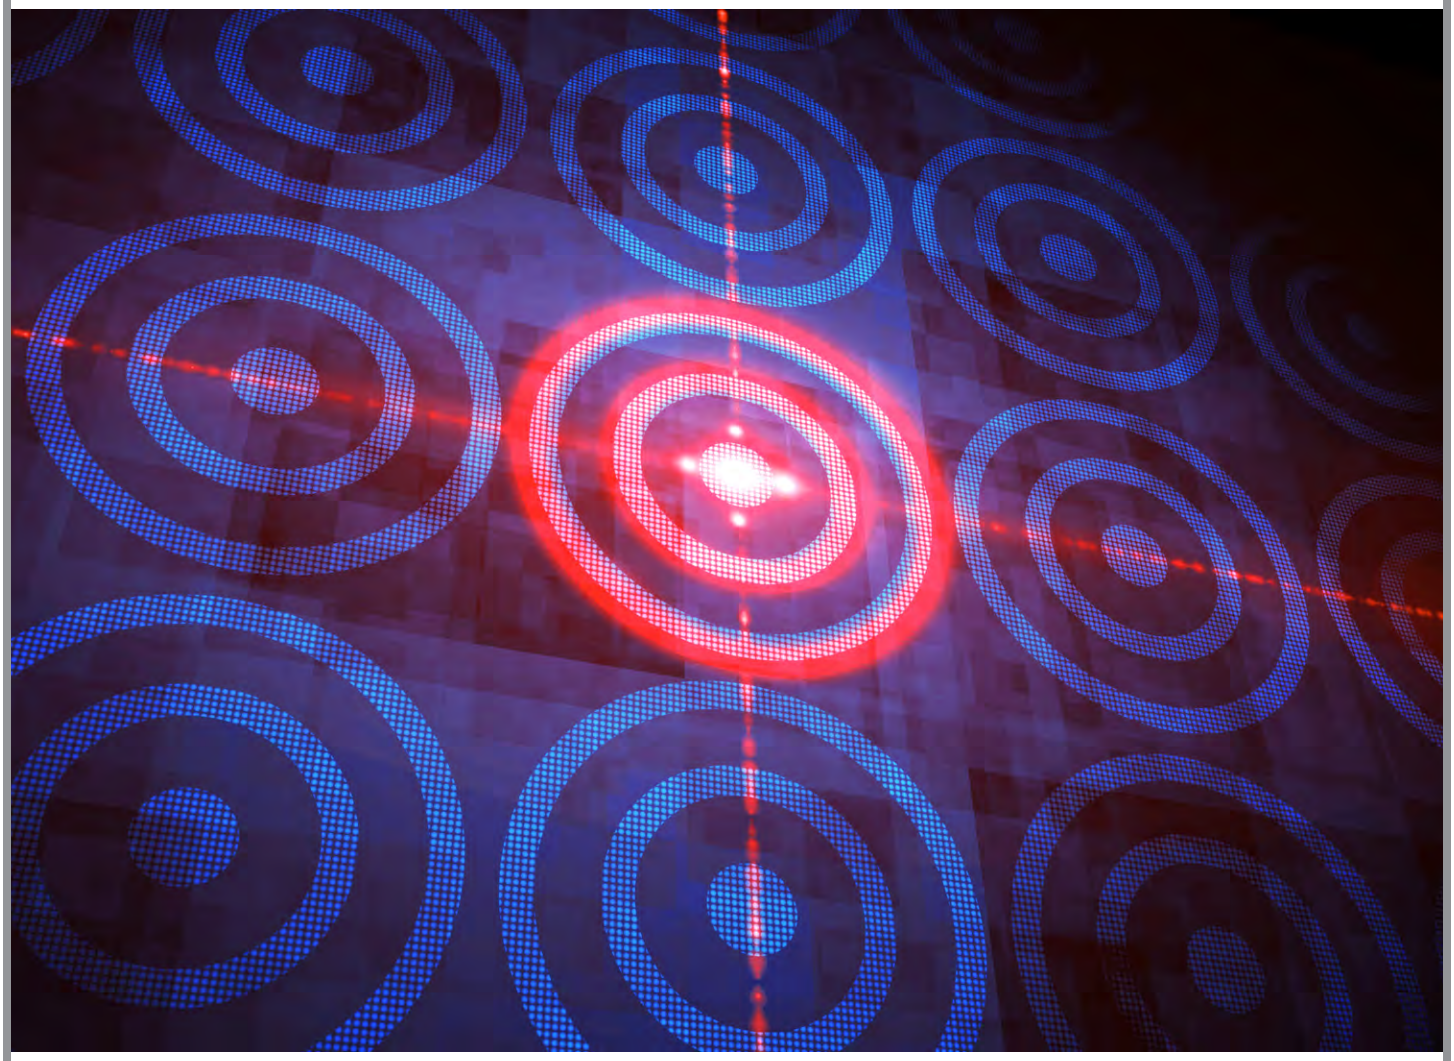

UNITRENDS

TABLE OF CONTENTS

UNITRENDS

Today's datacenter is evolving by means of several disruptive shifts in both methodology and technology.

For every organization or cloud provider, the datacenter represents every critical bit of information, application, and system you care about. So, naturally, those of you in charge of a datacenter continually look for new ways to protect the deliverability and availability of the data and applications you and your customers value most.

What happens, though, when the datacenter itself begins to evolve?

As the datacenter changes in ways you hadn't predicted, you need to continually update your data protection strategy to change with it. The last major change in the datacenter was virtualization. Your number of physical servers dropped, while your responsibility increased to keep even more applications running. You suddenly needed the ability to backup and recover these "virtual machines" (whatever those are…) and the environment they ran on.

But, virtualization was not only the "last big thing" to happen to the datacenter; it also was the foundation for several of the "next big things" to happen to the data center. All of which are causing you to, once again, shift your strategy to ensure you can hit the moving target that is the datacenter.

Today's datacenter is evolving by means of several disruptive shifts in both methodology and technology. It's these shifts that are making it more difficult for you to ensure you can protect the datacenter as it slowly turns into its next iteration. Let's take a look at each of these shifts and see how they impact the availability and recoverability of the next generation datacenter.

Shift #1: Agility

The first shift in the datacenter is the move to become more agile in its' operations. Neither IT, nor development, can function in silos any longer and must work together to adapt to the changing environment. Organizations wishing to be more agile realize they can no longer lock themselves into simply using virtualization and need to adapt to new technologies and methods that will provide improved performance, faster time of service, lower cost and better return on investment.

Virtualization gave birth to the automation of provisioning, maintenance, migrations and more. And with this automation (that we all now consider the norm), as the datacenter adapts, organizations must design and support environments that can automate not just the simple management of virtual environments, but far more advanced tasks such as self-healing, load balancing, and adaptive decommissioning.

As the technology changes or the needs of the business unit change, the datacenter needs to be agile enough to quickly more to new computing paradigms and not be locked into what will soon become the next "last big thing."

As the technology changes or the needs of the business unit change, the datacenter needs to be agile.

Protecting Through the Shift

Going agile in the datacenter will cause protection strategies to equally need to be agile in their approach. What will be required is more than just a willingness to ensure the protection of the datacenter; it will require keeping abreast of the changes in the industry that may impact your datacenter, so you can be ready the day that next shift occurs and you suddenly have one of those "next big things" sitting in your environment with the expectation of backing it up.

As you will see throughout the remainder of this whitepaper, protection will need to encompass new types of datasets that never existed, a resurgence of the use of physical servers and needing to think about protection of the datacenter by going outside the datacenter itself!

Shift #2: Containers

One of the newest developments in data centers is the concept of containers. Containers represent a more lightweight virtualization layer for applications to work within. So rather than hosting an OS within a hypervisor (and an application within the OS), containers serve as a viable alternative. Companies like Docker are leading the way with a container-based platform for distributed applications.

Running as a layer on top of the host OS, containers have some significant advantages over virtualization. First, containers allow applications to run with increased density (for example a single server might host 20 applications within containers instead of just 7 virtualized servers hosting

applications). Next, containers provide for improved application performance (which becomes somewhat obvious when you consider containers eliminate the need for the host OS to use resources to run the guest OS in addition to the application). Lastly, given the first two advantages, containers are good for an organizations budget. When capex is tight and needs to be stretched, being able to run 20 applications on the same server you used to host just 7 on is a good thing.

Protecting Through the Shift

The day you run your first application in a container, you'll obviously need to figure out what makes up a container in order to ensure it can be recovered. But expand your thinking beyond today's new shiny container, as it only represents the shift that's occurring today. Tomorrow's "container" will be something else new that you will need to add to your collection of assets to protect.

Bare metal cloud is the fast provisioning of a physical system with a customer's specified design.

In some cases, the shift won't be to something new and unfamiliar; the next big thing may just be something you've seen your entire career.

Shift #3: Bare Metal Cloud

In an ironic twist of fate (one that gives hope to those that keep thinking next year their favorite technology will make a comeback), we see the introduction of bare metal cloud. Simply put, bare metal cloud is the fast provisioning of a physical system with a customer's specified design.

Now, you might ask yourself "why would you even want that?" The answer lies in your desire to have the fastest and most efficient datacenter. Some applications - specifically those whose workload is made up of big data - like using physical hardware without sharing server resources or storage. By using dedicated hardware (again, to a very particular design spec), application performance is far more predictable, particularly when you're talking in the millions of operations.

So instead of trying to get that last few percentage points of performance out of the virtual infrastructure by implementing technologies that speed up processing and disk performance, the use of physical servers provides the edge (and the performance) that is necessary.

Rather than the "traditional" datacenter or cloud environment being made up of just virtualized servers hosted in a number of locations worldwide, organizations and providers have begun to weave hosted physical servers into their virtualized environment, designating that smaller set of physical servers to take on the higher visibility, tier 1 applications. So the irony here is that physical servers are now creating a homogeneous layer, of sorts, in addition to the one virtualization brought us years ago.

Protecting Through the Shift

This shift requires more than just a passing "OK, I'll backup the physical servers too." It requires first recognizing the importance of the datasets and processes occurring on bare metal cloud servers that require dedicated hardware. It then leads to discussions around separate service levels, recovery point and time objectives, and whether disaster recovery to a virtualized instance is even an acceptable option. And then the protection planning starts. This shift is pretty important and will require some of that agility brought up earlier in the whitepaper.

The last shift in the datacenter continues this idea of integrating something familiar into the operations of the datacenter itself.

The last shift in the datacenter continues this idea of integrating something familiar into the operations of the datacenter itself. In this case, though, the "something familiar" may not be something you ever though should be considered as part of the datacenter.

Shift #4: Software-Defined Everything

The next generation datacenter won't be only defined by what technologies are used inside; it also will be defined by what it interacts with outside its four walls. While this sounds a bit futuristic, it's perhaps not as far in the future as we might think in terms of how we get our applications to talk to each other and, ultimately, to the end users.

The idea of a datacenter that is software-defined means the networking, computing, storage and storage networks all become part of a virtualized computing infrastructure that is delivered as a service. Automation is key and is managed by intelligent software, rather than the hardware or people overseeing the infrastructure. (A great example of this occurring today is Cisco's Unified Computing System, where virtualization, servers, storage, infrastructure, and management are all combined into a single offering.)

Remember, the datacenter isn't just a place to crunch a lot of numbers; for some organizations, it is where their entire business takes place. For example, a Managed Service Provider may be using a datacenter to monitor it's thousands of customers and, as part of that monitoring, they need to notify both the customer of a pending hardware failure and dispatch a technician. It's not the email or a text message that is the example of Software-Defined Everything; it's the idea that the datacenter itself automatically notified the customer and dispatched the technician.

The datacenter is undergoing a similar shift, that has already started in a lot of ways, to not just interact with other servers, applications and pieces of hardware, but to utilize non-traditional sources of information, and provide that information to equally unconventional destinations. It's sometimes called the Internet of Things. No matter what buzzword you use, it still represents another shift in the modern datacenter to utilize systems, nodes, clients, sensors, and more – all external to the datacenter – to provide more than it does today. One could envision a day where sales reps using a modern-day version of Google Glass identify a customer, run a background check on their credit and approve a purchase – all without a human entering a single piece of information.

One could envision a day where sales reps using a modern-day version of Google Glass identify a customer, run a background check on their credit and approve a purchase – all without a human entering a single piece of information.

There's no telling how far the reach of this automated, software-defined datacenter may go, but you can be sure you will be asked to keep it running.

Protecting Through the Shift

This will be one of the more interesting shifts, as it's a complete unknown. The previous three shifts share elements of shifts we've seen in the past, but the very concept of delivering a datacenter that is undergoing major fundamental changes, building software-defined automation around it, and then delivering it as a service is a bit like changing the tire on your car while driving at 150 mph. Add to that automation that interacts with external devices that you never conceived should be involved in operations, and you begin thinking about the possibilities.

Can Data Protection Hit This Moving Target?

We're operating in some pretty interesting times. Datacenter operators and Cloud providers will be undergoing some significant changes that will require the rethinking of their protection strategies.

Protection and recovery needs to be thought of in terms of the many VMs, containers and bare metal cloud servers that are needed to recover an application.

Like the datacenter, data protection and recovery needs to not just adapt, but actually become adaptive in nature; constantly being able to adjust with the changes in the datacenter.

It's no longer a question of recovering your data, or an application or even an entire server; the issue data protection has is can it recover your entire IT infrastructure intelligently. Protection and recovery needs to be thought of in terms of the many VMs, containers and bare metal cloud servers that are needed to recover an application, taking into consideration the dependencies they each have and the order in which they are required to be recovered.

And lastly, protection (both backup and recovery) needs to be just as automated as backups, infrastructure management and anything else your software-defined datacenter is running to stay operational. Protection can no longer be though of as "that thing you do when you're notified a system went down"; it needs to be an integral part of operations.

Even in light of all the changes that are coming to a datacenter near you (and some already have), it's not only possible for datacenter protection to be successful, but is already happening. You're going to need to look beyond the datacenter you have today and begin to make some shifts of your own.

UNITRENDS