Protecting your Business and its Data: Backup vs. Disaster Recovery

BIOS

Be secured. Be assured

Contents

CloudHPT

Simplicity in the cloud

Security without compromise

Availability by default

Support without equal

BIOS Middle East and Veeam: The best choice to host Veeam backup and DR

BIOS Middle East contact details

About BIOS Middle East

Introduction

Are we back up and running yet? It's one of those dreaded questions no IT pro or business owner ever wants to ask or hear. It may come as a surprise to know that protecting your company's most valuable data, systems, and applications doesn't need to be in question, but instead can be something that's planned for and systematically executed when the time comes. So, instead of wondering if you'll be back up soon, you can be in a situation where you know *when* you'll be operational again.

It's one of the most overlooked services organizations should have in place. Some dismiss it as an overpriced insurance policy. Others see it as not being worth the effort, when you can just rebuild a system or application should it fail. But, nothing could be further from the truth. You don't just want to be in business, but to remain in business – and there are a lot of threats that can keep you from doing so.

So, why does your data need protecting in the first place?

There are a few common use cases for data protection for any business. Some apply to every business, while others are more specific to regulated industries...

- Data retention: Old versions of data, or information that no longer exists may be needed in the future to address an operational need or legal issue. Just like archiving, many organizations look at backups to strategically retain years of data without needing to maintain a lot of storage and expense.

- Disaster recovery: Whether the "disaster" is a deleted file, an office destroyed by natural disaster or anything in between, data protection is necessary as a proactive step to ensure your organization can recover to an operational state.

- Ransomware and cyber-attacks: Somewhat a variant of disaster recovery, many organizations are feeling the ever-increasing sophistication of ransomware that is responsible for the manipulation (via encryption) of data – making it useless to the organization without first paying a ransom for the decryption key. The only way to truly recover in these situations is to have backup copies of the impacted data, applications, and systems.

- Regulatory compliance: Some organizations are subject to compliance mandates that include specific data retention durations (e.g. seven years of audit data for SOX), safeguards against potential threats or hazards (e.g.. electronic patient health information in HIPAA and consumer financial information in GLBA), data integrity (e.g.. SEC/NASD) and data security (e.g. found in PCI, HIPAA, FISMA and more).. In regulated industries subject to these and other compliance standards, IT needs to prove they have a robust backup and recovery strategy to ensure data security and integrity are maintained in the recovered instance of the subject data.

- Customer expectation: With the digital transformation of businesses well underway (where traditional brick and mortar businesses are transforming into online operations), customer expectations of business Availability is quickly shifting from Monday through Friday to 24/7. Being able to keep your business running and available to customers is critical to staying in business today.

There are plenty of reasons data protection is necessary. But, as you try to build out a data protection strategy and you start thinking about actually using your backups at some point, you enter the world of backup buzzword overkill. You know the dreaded place where you're told you need backup, disaster recovery, recovery as a service, disaster recovery as a service – and you're left wondering which of these options is the best for your organization.

So, what is the right choice to best protect your organization's data?

What exactly is backup?

In the world of data protection, your backups are the copies of your virtual and physical systems, including all the data and applications in those systems. The goal of a backup is really, just to have a copy of the IT environment that makes up your business. In many ways, backups are more an insurance policy against something that may occur in the future – but with no definition of exactly what may come. While they certainly provide value – as it's far better to have backups than not – as you'll see, the methodologies behind backups help to address data retention and, marginally, some level of disaster recovery, but do little in the way of specifically addressing all four previously mentioned data protection use cases.

How often are backups created?

Backups are typically created daily, with no tiers of criticality or priority. It's like making a photo copy of everything on your desk at the end of each day, with no regard for what is and isn't important. Backup sets are preserved at varying intervals to establish a long-term data retention strategy. They often look something like the following:

Backup type	Retention
Daily	1 Month
Monthly	1 Year
Quarterly	> 1 Year

It should be noted that the *monthly* and *quarterly* backups are nothing more than a daily backup taken on the last day of the month or quarter. You'll note that nothing in this strategy specifically addresses a certain disaster scenario or compliance mandate, but as a retention strategy, these types of models are often mandated by operational or regulatory requirements.

Where are backups stored?

Backups can certainly be stored on-premises, but in recent years, the cloud has become an extremely viable storage option due to its geographic diversity from your main site, low cost storage options, and straightforward management. The general rule of thumb for backups is called the *3-2-1 Backup Rule*:

- Have at least three copies of your data (one is your production data)
- Store the copies on two different media (disk, tape, USB drive, etc.)
- Keep one backup copy offsite (which can be the cloud)

So, you can see backups as a data protection strategy provide some level of value. You have retention that allows you to recover back to days, months, or even years ago, with those backups existing in more than one location on more than one media to provide redundancy.

But what's missing in the conversation is the "what are we going to use this for" aspect to backups. Backups provide you an ability to recover, but aren't strategically designed for recovery in various disaster scenarios.

And that's where disaster recovery comes in.

What is disaster recovery?

There are two distinct differences that define disaster recovery (DR) from backup – and they are right in the name. First off, the focus of DR is to recover – DR is a system of replication designed to minimize downtime. It creates a copy of the VM at a secondary location and can fail-over in seconds or minutes. The second focus of DR is that it is *disaster*-oriented. To be clear, a disaster doesn't necessarily need to be a hurricane or flood; the disaster in DR is more about any kind of scenario that causes downtime. With DR, instead of simply having backups the intent is to devise a recoverability strategy for when a given disaster strikes – and be able to failover production systems and get the business back up and running very quickly.

How is DR different from backup?

There are a few factors that make DR different from backup alone. DR today usually involves some level of continuous replication or snapshot technology using image-based backups of either physical or virtual machines. These backups are used to bring systems back into an operational state, allowing a business to resume operations quickly. DR also utilizes two objectives to define what the recovery for a given system, application, or data set needs to look like. They are:

- Recovery time objective (RTO): While backups are defined by the data set to be copied, DR data sets are shaped by how long you have to recover. The RTO dictates the amount of time the recovery must take place within.
- Recovery Point Objective (RPO): When you recover, will it be back to an hour ago? Yesterday? Last week? The RPO dictates how much data can be "lost" when the system in question is recovered.

By putting these objectives in place, you are forced to create a backup methodology that will facilitate both objectives being met. This is normally accomplished on a per-system or per-application basis, to allow for recovery to align with the needs of the business. For example, should a critical application need to be up and running within 15 minutes and lose no more than 15 minutes of data, the RTO and RPO are both 15 minutes.. So, image-level backups become image-based replication, and a daily backup frequency becomes a differential replication of a virtual machine every five minutes.

Disaster recovery can also leverage a secondary on-premises site, or utilize Disaster Recovery as a Service (DRaaS).. DRaaS involves the engagement of a cloud service provider that facilitates some or all the recovery process and hosts the replicated systems in their cloud. DRaaS provides more benefits to the organization than secondary sites by providing geographic diversity (think along the lines of having a virtual secondary location that is anywhere in the world other than the area impacted by, say, a hurricane), lower costs without the secondary data center, and the support of an engaged third party to help in an emergency..

So, now that you know the difference between backup and disaster recovery, *is each of use, or is the answer just disaster recovery?*

Which one is right for you?

The answer isn't always *disaster recovery*, as certain scenarios can easily be addressed through frequent backups. Let's take the following two scenarios and see how each can be of use to your business.

Backups get it back

Because the backup data sets and frequencies are limited, backups are perfect for instances when alternate locations, third party DRaaS providers and replication are, well, overkill. Take the example of a deleted folder, or a workstation that has been disabled due to ransomware. Time sensitivity – from both a recovery point and recovery time perspective – isn't a factor here (this, putting aside the fact that you of course, want to get either of these lost parts of the business recovered in as short a time as possible).. Restoring either the entire system (as in the case of the workstation) or the deleted folder from the last backup prior to the issue at hand will meet the need here.

Disaster recovery (generally) puts it back somewhere else

In situations where the business is more concerned with an inability to recover on-premises (or perhaps do so quickly enough), disaster recovery is far more beneficial to you than backups. Take the example of an electrical short in the server room that causes the fire sprinklers to kick in. Every piece of IT equipment is damaged or ruined – and that may include your backups. At this point, the goal is to get your most critical services back up and running in far less time than days or weeks. Disaster recovery measures – such as continuous replication, where changes to a server are backed up and automatically applied to a mirrored system at an alternative location – will easily address the need, putting those services back into commission with little or no loss of data.

Choosing the right one

You might think it's a simple matter of "if it's a server, it needs DR." DR can be a costlier option when inappropriately used, so it's not something you necessarily choose across the board. There are a number of considerations you should apply to each workload (that is, the set of systems, services, and applications that make up part of your business – all of email, for example) to help you determine which methodology you should use.

- Workload criticality: This is a bit subjective, but how important is the given workload to the business? Can you survive without it? And for how long before the business will actually start suffering? In the world of data protection, this is called the *Maximum Tolerable Period of Disruption* (MTPoD). By establishing an MTPoD for each workload, you take the subjective "how important" question and instead objectively prioritize your workloads. At some point, you'll see the dividing point, where one set of workloads need to be up "now" and another that can wait a much longer time to be recovered. That first set is a prime candidate for DR.
- Rate of data change: The frequency of changes to a given workload usually align with its criticality in the organization. Take an order tracking system for a clothing company receiving hundreds of orders an hour. Should that application go down, can they afford to lose, say, an hour's worth of orders? Of course not. The RPO for an application like this is likely measured in single digit minutes. So, as you are parsing through your workloads, determining whether each should go the backup or DR route, high rates of change usually require DR.

- Cloud vs on-premises recovery: Note, this consideration isn't about where you should store your backups; both backup and DR can take advantage of the cloud for storage of backup data. This consideration is about where you will recover to. It should be noted that DR doesn't necessarily need to be cloud-based, but to maximize the ability for an organization to quickly recover operations based on the other considerations in this section, the cloud is an obvious choice for DR.

- Spend: There is a capital vs operating expense (capex vs. opex) discussion that also needs to be taken into consideration. Backups tend to be a capital expense, trying up budget for backup software and storage. Disaster Recovery, if done strictly with on-premises hardware, can also drain IT budget. But it's the use of cloud-based backup storage and DRaaS services that can shift the discussion from capital to operating expense. With little more than making a monthly payment, you employ the same levels of data protection, recoverability and Availability as much larger organizations – without needing to invest heavily in hardware and software.

Determining which option is right for your organization requires walking through each of these considerations on a per system/application/data set basis. There is no wrong answer here – the right answer is one that meets your business expectations and needs. Involve the C-suite, application owners, line of business owners, IT, and even more technical users to walk through the considerations, applying them to the critical workloads within your organization, to identify whether backup or DR is the right choice..

Regardless of which option is chosen, you need to make certain the hardware, software and/or service provider assisting in backing up and recovering your operations can meet your specific data protection needs.

Getting the most of your data protection provider

You don't just want to choose, say, DR and then pick a vendor solely because they state "we do disaster recovery!" You need to look more diligently into the services they provide, how those services are provided, and whether they meet your recovery needs and budget constraints.

So how do you maximize either your backup or DR efforts to come up with the most cost-effective overall solution to meet your recovery needs?

The answer lies in partnering with the right backup or DR provider, leveraging their experience and expertise to craft your data protection strategy. Use the following sets of questions of each type of data protection vendor to customize your strategy and identify the right vendor for you.

Questions to ask your backup provider

Most of the questions you need to ask revolve around what should be backed up, how often and where it will be stored. Additionally, there should be some questions asked around how to best access and utilize any backup data stored in the cloud.

1. Can you help with prioritizing workloads? Rather than just having your organization subjectively prioritize workloads, having a partner who understands what's possible from a backup and recovery perspective will help to properly put your workloads into order.

2. How frequently should backups be configured to run for each workload? Daily isn't necessarily the answer. Some will need to be backed up more or less frequently. Understanding the capabilities of your backup solution with regards to the workload being backed up will help set expectations.

3. How much storage will I need? Depending on the backup and retention strategy, this answer will vary. Answering the first two questions will assist in coming up with a pretty close estimate.

4. Where should my backups be stored? You have a few options here; on-premises, offsite copies of physical medium, and the cloud. The answer may vary and will be dependent upon the criticality of a given workload.

5. When are backups to be saved to the cloud? Should all or part of your backups be stored in the cloud, you need an understanding of how the backup data is stored in the cloud. Is it first stored locally and then replicated? Sent directly to the cloud? How quickly?

6. How do I retrieve backups in the event of a disaster? With the assumption that you'll be performing recovery operations yourself, you need a full understanding of exactly how to get your data back. With most cloud-enabled backup solutions, the recovery from cloud to on-premises is seamless.

8. Are there additional "hidden" fees? Be certain of the services being purchased; you may be thinking you are getting more "aaS" (where the vendor does the work in your time of need) when all you're contracting for is purely the backup services and infrastructure. Some vendors also charge for setup, the number of systems/VMs being backed up, and even bandwidth used when backups are stored in the cloud. Know what it's really going to cost you – and what you're getting for your money.

Questions to ask your DR provider

The goal of your questioning should be to better understand how they plan on establishing an environment that will facilitate meeting your recovery objectives. Some of the questions below (particularly those around the security of your data) also can apply to backups.

Note: Many of these questions assume you will be replicating image-based backups to a cloud-based virtual environment hosted by a DR vendor.

1. What types of network configurations can you accommodate? Your organization's most critical applications may not just be on-premises. Understanding whether the DR vendor can address protecting both on-premises and cloud-based applications is important.

2. Is the target hardware enterprise-class? DR vendors should have the same hardware ready and waiting for *every* customer – regardless of the customer size, data size, or DR need. Hardware should be able to duplicate your environment down to the last detail, and provide equivalent

(or better) levels of redundancy with regard to accessibility and Availability to your organization.

3. What is the impact on my on-premises systems? Like backups, this question needs to be answered to ensure the backup efforts at an image level have as little impact on production systems as possible. Additionally, in a situation where only a part of your environment is failed over to a DR site in the cloud, you should know whether your still-running systems will interact as normal with recovered systems residing in the cloud.

4. What types of systems do you support – virtual or physical? You need to be sure your hypervisor is supported. Some applications perform better (particularly data-intensive apps) on physical hardware. So, your DR'd version may need to also be a physical system. Knowing whether your DR vendor supports both your physical and virtual systems – as well as whether they can recover to their equivalents in the cloud – is crucial to ensure continued productivity.

5. How much bandwidth should be available for replication? This is solely based on the amount of data you must back up, the frequency and amount of changes, and the frequency of backups to achieve recoverability.. A good DR vendor (once they've helped define the backups needed for recovery) should be able to tell you what to expect.

6. Who will help, in the event of an emergency? You are looking to a DR vendor to help you in your time of need and knowing just who will help, but the plan when it happens is something you need to know proactively. Do you call someone specific? Will they already know of a failure? Get the specifics and be sure you're comfortable with the answers.

7. What kind of support do you provide during a DR event? Beyond just contacting your DR vendor, you need to know what parts of the DR process they will be assisting with. Are they merely storing the backups and providing you with the recovery infrastructure? Are they performing the recovery for you? You should also know the kinds of support options available – email, chat, phone – and the service level agreement (SLA) that is part of your contract. Do you have 8x5 support? 24x7? Will they help if you have an OS or application issue during recovery? Be sure this is all documented in your SLA.

8. How fast can the systems fail over? This is both a question about the DR capabilities of the DR vendor, as well as a question that goes back to whether you have the right backup strategy in place. For example, the DR vendor may tell you they can have a system back up within 30 minutes. Not bad – that is, unless the organization needs that system up in 10, in which case you need to re-evaluate your backup methods to ensure a faster recovery.

9. Are you providing partial or full failover? This is all about whether the DR vendor is failing over all of your environment (or all systems related to a given application) or only a part of it. If it is a partial failover, you need to be sure your intact systems will work properly with failed over systems. You should also plan on testing a failover of critical apps if only partial is supported. If a full failover, dig a little deeper and inquire about how the failover is performed – are there dependencies established so certain systems fail over before others? Will automation be in use? All this should be defined for you.

10. Where will my systems live, in the cloud? At the end of the day, if the business can continue to access and use the failed over systems with little or no interruption, you may not be concerned. However, you should be aware of where in the world the data resides, what kind of Availability guarantee you have while your systems reside there, whether it meets your geographic and data sovereignty requirements, and if it meets the compliance regulations you are subject to.

11. How secure are my backups? You're potentially putting your most sensitive and critical data into someone else's hands. How is the data secured in transit? At rest? What levels of encryption are used? How is the recovery environment protected from external attack? Does the environment meet the ISO 27001 information security management certification? Be certain to have a grasp on exactly how secure your data is, both when stored as backups and when used for recovery.

12. How is the service priced? Pricing can vary between vendors to include monthly access fees to reserve recovery infrastructure, usage-only-based charges once systems are failed over, or a mix of both. Additionally, costs for services like monitoring, failover, and failback can also be offered to you.

Whether you are going the backup or DR route, it's critically important to have a full understanding of what you can expect – in both functionality and service – from your hardware, software, cloud storage and "aaS" vendors.. Know how they will help you plan and prepare, how, when, and where backups will be generated, what recovery will look like and who is responsible for it.

Putting data protection into practice

Of everything presented in this ebook, what's most important is *that you go through the exercise* of looking at your workloads, breaking them up into tiers of criticality, and determining what's the best way to ensure they are protected and made available to keep your business running.. As you do so, there will be obvious opportunities each for backup and DR, and workloads that may sit somewhere in the middle. For every workload, find a means to protect it, and put that protection into practice using a partner who specializes in data protection.

Veeam backup – how it works

Nearly every organization today has most of its critical workloads virtualized. With an ability to quickly move, revert, and replicate virtual OS images, virtualization is a natural choice that adds flexibility and performance to an Availability strategy. Veeam®'s backup efforts center around protecting both VMware vSphere and Microsoft Hyper-V virtual environments. Backing up at an image-level, Veeam performs application-aware backups of entire systems, ensuring application consistency and recoverability. VMs can be automatically tested and verified proactively to ensure recoverability before it matters.

Keep it on-premises

Backups can be stored locally as a first line of defense, integrating backup acceleration technologies with tape and storage vendors to generate extremely fast backups used to meet tight RPOs at time of recovery. Backups can be generated from storage snapshots when using storage from vendors like HPE, NetApp, EMC and Nimble, as well as via direct storage SAN and NFS access. For larger environments, virtual pools of backup storage can be utilized to allow for scaling storage capacity without reconfiguring backup storage targets.

Leverage the cloud

Utilizing Veeam Cloud Connect, backups can natively be stored in, or backed up directly to the cloud, using cloud-based backup service providers for additional layers of recoverability. Backups use built-in WAN acceleration to ensure the smallest amount of bandwidth is used during backup and recovery operations.

Simplify recovery

Veeam provides the ability for reliable, flexible recovery at a VM, file, and application level. VM recovery is an easy process that takes just a few minutes. Files and folders can be recovered just as quickly. With Veeam Explorer™ applications, recovery of objects within applications and platforms like Active Directory, Exchange, SharePoint, SQL Server, and Oracle are lightning fast. Recovering a user account, an email, or a database table are as easy as navigating within the appropriate Veeam Explorer application and selecting the object in question to recover it.

Veeam disaster recovery – how it works

While fast recovery is definitely a focus for Veeam backups, DR with Veeam concerns itself with expanding recovery to complex distributed environments containing virtual and physical systems – all in an effort to encompass the recovery of all applications and data, while meeting SLAs for RTOs and RPOs of less than 15 minutes each.

DR for all of your environment

Veeam's backups for VMs integrate within the virtual environment itself. Leveraging agents specifically designed for Windows and Linux systems, DR efforts can be extended to physical machines to ensure every part of your environment is protected.

Complete visibility

Veeam provides real-time monitoring of backup and VM performance, alerting you to potential issues that may impact SLAs. Optimization of virtual resources, along with capacity planning and forecasting means you can maximize your environment's ability to not just run, but run at peak performance.

Data loss avoidance

It's critical for your backups to be complete, and without any missing data. Veeam uses application-aware, image-based backups to ensure the most up-to-date consistent backups without requiring system shutdowns. Backups can be replicated on site for higher levels of Availability, or to the cloud for use in disasters where the on-premises network may no longer be available.

Verified recoverability

Everyone plans for recovery, but rarely do you find testing as part of the process. And without a valid backup, your recovery is doomed to fail. Veeam employs SureBackup and SureReplica to validate your recoverability. SureBackup automatically boots up backed up VMs directly from the backup file, performing a test to validate its' status, and it emails you a state report. SureReplica tests restore points in VM replicas for consistency and reliability by running the VM to the associated restore point in a virtual lab.

High-speed recovery

Achieving a recovery time of less than 15 minutes, given the variety of recovery scenarios that can exist, can seem a daunting task. Veeam provides a number of means to ensure recovery – no matter the disaster. DR with Veeam can be accomplished at the same VM, file-level, or application object restores found with its backup capabilities. In addition, you can perform universal application item recovery without agents or empower users with one-click self-service VM and file recovery.

CloudHPT

Veeam's robust backup and disaster recovery technologies are only strengthened with the right partner that has created a cloud-based offering that seamlessly integrates with Veeam. Finding a cloud provider with this integration will add value to your backup and DR efforts in the areas of simplicity, security, Availability, and support. BIOS Middle East, the 2017 Veeam Service Provider of the Year, has over 10 years of experience in helping customers across many industries design, implement, and manage reliable cloud-based backup and disaster recovery strategies.

Simplicity in the cloud

As a Veeam Cloud & Service Provider (VCSP) partner, BIOS Middle East effortlessly integrates with Veeam backup and DR solutions to provide you with both easy backup and simplified DRaaS in the cloud. Whether using Cloud Backup with Veeam Cloud Connect or Veeam replication technologies for disaster recovery, the BIOS Middle East solutions can be configured and ready to use in only a matter of minutes. Customers needing both legacy and physical systems alongside their VMs can count on BIOS Middle East to protect every part of your environment that is needed as part of your disaster recovery strategy. And, with straightforward monthly pricing for both cloud backup and DRaaS, there are no hidden setup, per-VM or bandwidth fees to worry about.

Security without compromise

Providing end-to-end encryption, your data remains secure at the source, in-flight and at rest when in the BIOS Middle East cloud. Organizations governed by data security standards found within relevant compliance mandates can choose from data centers globally that meet ISO/IEC 27001:2013 - information security standard, ISO 9001:2015 - Quality management systems, ISO 22301 Societal security -- Business continuity management systems and ISO 20000-1: Service management system, which exceed the security standards found in most regulations today. With BIOS Middle East data centers located in the UAE, those customers concerned about data sovereignty can host their data with assurance that data will not cross borders.

Availability by default

If you rely on a cloud provider to help maintain your Availability, they must be available more than you. With a 100% infrastructure Availability SLA, it doesn't get any more available than this. Using state of the art best-of-breed hardware, strong multi-layer embedded network security measures, multiple redundant direct access connections, along with physical and environmental controls, BIOS Middle East's maintains the uptime necessary to meet any backup or DR need – whenever and wherever it may happen.

Support without equal

Despite best efforts, you still may have recovery problems during your time of need. BIOS Middle East does more than just provide you with the backup and recovery infrastructure; they include 24/7 phone and email support by certified experts to assist issues directly related to backup and recovery, as well as provide access to system administrators, network engineers, security experts, and senior architects skilled in diverse infrastructure environments to help with getting you operational again.

BIOS Middle East and Veeam: a total Availability solution for your business

Protecting your company's IT systems and data is probably the most important task you'll ever perform. You need verified backups, redundant copies on-premises and in the cloud, a reliable cloud infrastructure to recover to, and the confidence this will all come together when "the data hits the fan." So, your recovery efforts – whether as part of backup or DR – need to employ a robust backup and DR solution and a provider with years of recovery expertise to ensure you have the solution, support and cloud platform you need to quickly, securely and consistently back up and recover the data, applications and systems that define your business..

BIOS Middle East contact details:

www.biosme.com

800 BIOSME ((246763)
nidhi@biosme.com

About BIOS Middle East

BIOS is an award-winning Managed Service & Cloud Provider established for more than 16 years in the UAE and committed to helping organizations become more efficient and cloud enabled through in-country and global cloud networks and platforms. We provide over 400 customers with over 10,000 servers and the capability they need to achieve successful transformation, through industry leading, end-to-end managed networks, security and platform solutions

The core benefits of Managed Services & Cloud are well documented:

- Improved productivity and efficiency
- Elimination of Capital costs
- Reduced Operational Costs
- Dramatically Improved flexibility and business agility
- Ease of Compliance
- Enhanced Security posture

The challenge is implementing a Managed Service and Cloud strategy successfully so that these benefits are actually achieved. In many cases, the wrong decision or approach can lead to increased costs and headaches. BIOS offers a simple solution to this problem. We have all the components required to help a business to manage on premise datacenters and or transition to the cloud, from a team of expert consultants to our own in-country cloud, and everything in between.

We are vendor, platform and technology agnostic (we are a fully certified and accredited partner with all the major players), ensuring that the solutions we provide are tailored to suit your exact business requirements. Working with you at every stage of your cloud journey – from discovery to delivery – we'll help you navigate the complexities involved in building and operating cloud systems.