WHITEPAPER

Security benefits of a good cloud service

How to plan and configure your cloud service to maximise security and business benefits.

Introduction

At the NCSC, we often hear from customers who are concerned about the new risks they associate with moving their traditional systems to the cloud.

For example, there's a perception that moving to the cloud means having to invest much more trust in your supplier. At the same time, the idea of moving to a different operating model can be intimidating. There's a belief that moving to the cloud makes it easier for organisations to accidentally provide the public with access to confidential data.

Given these perceived risks, the security benefits of moving to a cloud-based system are often overlooked, with some organisations believing that cloud deployments will result in more overall risk than traditional systems.

This security paper seeks to challenge these perceptions. It describes the security benefits from adopting a good cloud service. We'll describe a number of security outcomes that have proved difficult to achieve in traditional IT deployments, and how these will benefit your organisation. It's important to realise that these benefits won't be fully realised if organisations adopt a 'lift and shift' approach; the cloud service will need to be planned and configured to take advantage of the benefits that a cloud-first approach can bring.

> **Note on audience:**
> This paper is aimed at risk owners considering a move from on-premises to cloud-based solutions. Technical staff building digital services in the cloud may also find it helpful. Note that while this paper is focussed on public cloud, many of the benefits can also apply to a good private or community cloud.

Terminology

A few quick points on terminology before we start.

Automation

Any cloud service will make extensive use of automation, which means a task will be performed by software, to ensure the task is performed consistently. This may

be anything from a small script to a large, distributed application. It's important to note that the task may still be initiated by a human, so the automation is not necessarily autonomous, but the bulk of the task will be conducted by software.

Availability zone (AZ)

Cloud providers typically use terms like 'availability zone' or 'zone' to refer to one or more data centres. The data centres in an AZ will normally share resources, such as power or cooling suppliers, connectivity to the internet, or physical location. In contrast, two different AZs should use independent suppliers and be geographically distant. This prevents most issues from affecting more than one AZ at a time.

Cloud

When we use the word 'cloud', we're referring to the NIST definition which describes it as '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'.

Cloud provider

For the purpose of this paper, we consider the vendor and operator of the cloud services you use to be a single entity, which we refer to as the **cloud provider**.

Cloud service

Moving to the cloud tends to mean using any number of cloud services. These vary substantially in size, from an entire e-business suite to a single component within a software development ecosystem (such as a storage or cryptographic key management service).

Guardrails

An emerging technique in cloud is the use of automated mechanisms to enforce security or policy requirements, called guardrails. This could involve a requirement that all bulk data storage encrypts data at rest, or that only a limited set of approved operating system images can be used for compute service.

Infrastructure-as-code

Cloud infrastructure, from more tangible things like VMs to the configuration of more abstract services, is normally managed using APIs. This makes it easy for the cloud service to take a text template that describes the desired configuration and call the APIs on your behalf to make the necessary changes. This means you can track your configuration in text documents, where changes can be reviewed and automated analysis can be performed.

Outcome 1: Self-service with guardrails

Traditional IT often requires a member of staff to authorise new equipment and resources, such as virtual machines, firewalls, or internet-facing IP addresses. This kind of authorisation tends to be slow. If you've ever had to submit a change request to start a new VM, you'll be familiar with this pain. However, these checks serve an important purpose, by preventing undesirable changes. For example, it's good practice to have an inventory of internet-facing services, so that the security operations centre (SOC) can triage security alerts about those services more effectively.

In the cloud, the same requests can be made self-service, with automation performing the checks and balances previously done by staff. Automated guardrails can be configured to ensure undesirable changes are blocked, without imposing unnecessary delays on acceptable requests. For example, guardrails can ensure data storage buckets are encrypted and private by default and ensure logs are fed into the central aggregator. This approach balances the need for rapid change (including responding to security issues) while preventing undesirable changes. At the same time, having these sorts of checks performed by automation can give strong confidence that the system is behaving how you expect.

How to maximise the benefit

When configuring or building guardrails, it's important to understand what checks would previously have been performed on the request in question. It's also important to identify any new checks that should be considered. For example, is the resource tagged with owner and cost centre information? If the guardrails

reject a request, do they give a clear, meaningful, and actionable error message? Can a guardrail prevent an undesirable change or just fix the problem afterwards?

Outcome 2: Automated service management

Secure service administration focusses on tackling two key issues.

1. Reducing the likelihood that an external attacker gains control over the service's administration interfaces.
2. Reducing the impact of a successful attack by limiting the resulting access to the service.

This means putting in place strong defences around administration interfaces and strictly enforcing the principle of least privilege. It's important to consider secure service administration in the context of both the provider's administration of the service as a whole, and the customer's administration of their use of the service. In the NCSC's Cloud Security Principles, these are principles 12 (secure service administration) and 9 (secure user management), respectively.

A good cloud service will give customers a way to automate administrative activities. For example, the service may expose APIs corresponding to manual administrative activities (like configuring access control or creating resources). This gives customer admins full control over what actions can be taken whilst administering the service. Using these automated approaches also means that a detailed audit trail is kept, tracking what has been done, when, and by whom. This is an effective deterrent against malicious insiders, as it cannot be modified or deleted by the customer or the insider. This automation also makes configuration management less prone to error, improving the customer's ability to use the service.

Similarly, a good cloud service will allow all users (including admins) to have granular permissions over the service. Traditional systems often treat admin users as 'all-powerful'. In the cloud, even admins should only have the permissions they need. For example, an admin might have full control over a limited set of resources or will have limited abilities (such as the ability to create but not delete

resources). If even admins have limited power, the impact to a successful attack on an administration interface is reduced.

While it's possible to put the same controls in place in a traditional system, it normally means investing a huge amount of effort for a benefit that can be hard to articulate. For example, most service administration today uses protocols like RDP and SSH, which tend to have strong authentication but don't produce a useful audit trail. Are you confident that you have an exhaustive audit trail of all the actions of your administrators today? Are you confident that the administrators have no way to modify the audit trail?

How to maximise the benefit

Make the most of this capability by using the service's APIs and other audited functions to perform administration on your use of the service. For example, rather than managing VMs using RDP or SSH, use a managed administration function if one exists. This will give you a much more robust audit trail of your administration activity. Make sure the audit information is protected and backed up securely.

Outcome 3: Programmable orchestration

Responding to incidents is challenging. Inevitably they happen at the most inconvenient time (such as when the team responsible is unavailable, or just after a change that means the prepared response in your run book is now incorrect). This all makes reliably providing a correct and timely response to an incident very hard. Customer use of cloud services is driven by documented APIs, which makes it possible to coordinate, monitor, and maintain a cloud environment autonomously. This enables customers to respond to incidents without requiring human involvement.

For example, common mistakes can be rectified by autonomous action. Taking this approach increases both the speed and predictability of the response. It also enables template-driven infrastructure-as-code, enabling reliable and repeatable deployment of services. Automation can also be used to create duplicate infrastructure for incident response and disaster recovery processes. Routine and

frequent tests using this approach gives confidence that the processes work predictably.

How to maximise the benefit

The most effective way to capitalise on programmable orchestration is to define your use of the cloud with infrastructure-as-code templates. For SaaS services, this will be the configuration of the service. For cloud platforms, this will probably include the structure of your applications, as well as the configuration of the serverless components you use.

If possible, use a continuous integration/continuous delivery (CI/CD) pipeline to deploy your templates. This will give you confidence that the templates behave as expected and that manual changes aren't being made on the side. You should also have a regular process for testing the templates from scratch to prepare for a disaster recovery scenario.

Outcome 4: Consistent test environments

Traditionally, each service built by a customer will have a small number of fixed environments, as changing the size of an existing environment (or adding a new one) is often costly and time-consuming. In a good cloud service, it's generally quick and cheap to build a new environment (and equally quick to delete one when it's no longer needed). Being able to quickly duplicate environments has many advantages. For example, changes can be tested on a duplicate environment, so that unexpected issues do not affect the live environment. If this testing environment is entirely built from infrastructure-as-code, then it will be identical to the live environment it represents. Having multiple environments also helps to prevent the compromise of one from affecting another.

How to maximise the benefit

Make sure you use infrastructure-as-code wherever possible, to make it easy to create duplicate test environments. Use duplicate environments to test your disaster recovery process in practice, regularly. Deploy your environments using a CI/CD pipeline to ensure that manual changes cannot be made to environments.

This helps to prevent the environments from ‘drifting’ from the template configuration.

Outcome 5: Encouraging good security

There are various ways a good cloud service encourages good security, both by automating a lot of the security for customers and by making it easier and more affordable through economies of scale. For example, by making small chunks of compute both highly available and cheap, tiered architectures can become more cost-efficient than monolithic designs. Limiting the duration of serverless functions (small compute components that perform a single task, such as data processing or 'gluing' two services together) makes them more cost-efficient for the cloud provider while also making them much easier to patch and harder to attack.

At the same time, expensive security components (like pervasive access control mechanisms and hardware cryptographic security modules) become affordable at cloud scale, where each customer pays a small charge for a slice of its use (rather than having to buy, maintain, and operate the entire component). This makes advanced security accessible to all customers, including those with very limited budgets. Additionally, the cloud provider's scale enables them to build up world class security expertise, which they can then bring to bear across all of their customers.

How to maximise the benefit

Make the most of the services available to you. The more high-level and managed a cloud service, the more cost effective it will be. Rather than reinventing the wheel in IaaS VMs, use managed services, which will probably solve the problem more effectively and more cheaply. This is particularly true for common tasks like serving static websites, which can often be free (or near enough) when using managed services, but quite expensive when implemented from scratch in VMs. Managed services will typically have a much more secure default configuration than building it yourself.

Outcome 6: Pervasive and consistent identity-based access control

A good cloud environment should have a single mechanism for authentication and authorisation. This will often use designs like role-based access control (RBAC), where groups of related permissions are stored in 'roles'. These roles can then be applied to users and workloads to give them the permissions defined in the role. This gives a consistent, familiar, and easily audited access control throughout the cloud environment, including for automated workloads.

Taking this approach means that granular permissions can be applied to each identity in the environment, enabling zero trust architectures and defence in depth against compromise. Secure access control is crucial to effective security and having a single, consistent approach dramatically reduces the chance of error.

A good provider will maintain an extensive set of pre-defined roles with clear naming and purpose, preventing most customers from having to define their own roles, which can be error prone. Authenticating each workload also enables secure key management and correct access control for core services like logging and monitoring. For example, most workloads can be given the ability to submit logs to the log aggregator service, but cannot overwrite or delete logs. Similarly, a workload might be allowed to encrypt new data but not decrypt old data.

How to maximise the benefit

Authentication and authorisation are two of the most important concepts in security. In cloud services, setting up your users' identities and applying access control is crucial to keeping your data secure. Spend some time to get familiar with how access control works in the cloud services you use. Prioritise simplicity when you apply access control. It's better to use the access control system in a way that gives slightly more permissions than you need, if doing so simplifies the design. Put in place monitoring to alert on the use of the most powerful permissions. Look for tooling that removes permissions that aren't used (so that you can maintain the principle of least privilege over time).

Outcome 7: Commodity expertise

Managed services tend to be configured, designed, operated, and monitored more effectively than any customer designing their own service could achieve. This includes solving hard security challenges (such as data encryption, access control, and integrated monitoring). This is because managed services can normally bring more specialist expertise to focus on a single problem. When customers try to tackle the same problem, it's just one small aspect of a much bigger system and one job amongst a huge list of other tasks.

By using managed services, particularly in cloud services, the provider's expertise can be used as a commodity. For example, a managed database service benefits from the vendor's investments in expertise in the underlying database product. This means that each customer benefits from more expertise than they would get if they had to buy-in this expertise.

How to maximise the benefit

Try to break down your system into smaller, less tightly coupled components. This should help you to replace many individual components with shared services. For example, if you can make one component responsible for the storage of bulk data, it can use a managed database or a blob storage service. You should be prioritising your effort on the aspects of the design bespoke to you. Rely on shared services for more generic considerations, such as monitoring, network management, and operating system management. Take full advantage of the cloud services at your disposal.

Outcome 8: Pervasive logging, monitoring, and audit trail

A good cloud environment will have integrated collection and storage of system logs, audit data, and monitoring metrics. This means that all logs produced by the cloud service (and most logs produced by customer workloads) will automatically be aggregated into a single location. In traditional networks, it can take considerable effort to aggregate all logs together, making it hard to gain a clear picture of what happens on the network. At the same time, many cloud services have integrated tooling for analysing and querying logs, making the aggregated data even more valuable.

A good cloud service will provide no way for a customer to interfere with the cloud service's logging, which makes:

- it harder for an attacker to cover up their tracks
- incident discovery and response more effective

While the visibility the cloud provider has within a customer-managed VM may be limited, this visibility improves dramatically in services managed by the provider, such as serverless and access control services.

As a cloud provider will often have more context around a service's behaviour than the customer, a good provider will add that context to logs and monitoring to make it easier for the customer to generate effective alerts. This includes access to sources unavailable to the customer, such as the underlying software-defined network and infrastructure. This additional monitoring also helps to protect customers from attacks their peers have suffered. For example, the provider can identify attack patterns suffered by one set of customers and mitigate issues before they have affected any other customers, like vaccine development against pathogens.

How to maximise the benefit

The most important first step is to ensure that all useful logs from your workloads are ingested into the central logging service. Then make sure that your log processing acts on the central log repository, so you can take full advantage of the logs produced by the cloud service itself.

Make sure you have strong access control on the central log repository, apply an appropriate retention mechanism, and back up any important logs. Perform regular drills where you investigate a set of activities using the logs and monitoring tools available to you so that your operations personnel are familiar with the tools and know what information is available.

Outcome 9: Abstracted network management

Network flow controls are a core mechanism for enforcing security requirements. Traditionally, network controls are applied using IP addresses and subnets. This works, but it can be quite prone to error. If the DHCP server has a bad day and a piece of software is given a different IP address, that can change what access it has. If an attacker can convince a machine to proxy its traffic, that may give them further access.

Cloud services can bring network flow controls to the application layer. Rather than dealing in IP addresses and subnets, customers can specify allowed data flows between applications. This makes network design easier to understand, less prone to error, and even introduces new security opportunities. For example, application layer networks may be able to add encryption in transit automatically and give more context to your network monitoring.

How to maximise the benefit

Look for services that take care of network management for you and apply authentication and authorisation automatically. For example, serverless components typically use the same access control system as the rest of the cloud for talking to one another. You may also be able to use services that apply application layer network flow controls on top of more traditional cloud networks.

Outcome 10: Effective patching

Prompt deployment of security fixes (also called patching) remains the single most important thing you can do to secure your technology. Some of the most severe attacks in recent history (such as Equifax's breach in 2017) would have been prevented if their systems had been patched effectively, although we recognise that patching effectively can be really hard. The pervasive automation used to manage a cloud environment means that patches can be tested and rolled out quickly at scale. Duplicate environments make patching easier to apply and help to identify problematic patches without damaging important systems.

At the same time, the relative uniformity of a cloud environment's hardware estate means that fewer patches will be necessary than for most traditional data centres. The incredible scale at which a cloud operates makes automated

patching of the entire stack a necessity, including the CPU, UEFI, and firmware of both servers and networking equipment. This greatly increases the likelihood that all components handling customer data will be fully patched. You might be confident that your OSs are fully patched in your data centres, but do you have the same confidence in the patching of your firmware and microcode, or the networking equipment you use?

Cloud providers often participate in vulnerability embargoes, meaning that cloud customers may often receive security fixes before major vulnerabilities are disclosed publicly. For example, the biggest cloud providers had already applied patches for the Spectre and Meltdown vulnerabilities before they were disclosed publicly. This puts cloud customers ahead of even the most diligent local sysadmin.

Finally, there will be some cases where an attack is published before a patch is available. Even in these cases, the cloud provider can often help prevent the attack, such as by blocking attack traffic in the software-defined network, or changing the separation between customers. While it's possible to use similar measures in traditional deployments, many organisations don't have the right combination of security and engineering expertise to apply mitigations on short notice. A good cloud provider will have the people resources and the expertise in all layers of their stack, making changes more quickly and easily.

How to maximise the benefit

The main way to benefit from a cloud provider's patching is to make them responsible for patching more of the stack. If you use IaaS VMs, they're responsible for patching the firmware and hardware, but the rest is up to you. With a managed container service, they take on patching the underlying OS kernel. With a serverless compute service, they take on patching the rest of the OS and possibly the programming language runtime. With SaaS, they patch everything!

Outcome 11: Authoritative inventory

Traditional data centres are typically billed by the space occupied and power capacity made available. By contrast, cloud services are billed by consumption of

each service used. This means that the provider must keep detailed records of all services used by each customer to be able to produce and justify the customer's bill. A good cloud service will use this capability to provide the customer a detailed and authoritative inventory of the services being used, along with each service's configuration and usage patterns. Having this kind of data makes incident detection and asset management far more effective.

How to maximise the benefit

Make sure you identify where your cloud provider makes the inventory available to you, and how much information it includes. Put monitoring in place to alert when unexpected resources are created to help detect suspicious activity. For example, if you're using a cloud service primarily for data storage and lots of VMs are created suddenly, that might suggest a compromise.

Outcome 12: Resilience

Improving resilience is a common reason for moving to the cloud. Both the distributed nature and the scale of a cloud environment means that the inevitable outages that affect all technology systems are less likely to result in user-visible availability issues. Cloud services can tackle outages by moving workloads to unaffected systems. Those systems should have the capacity to take on load from other areas, which reduces the likelihood of a cascading failure. Furthermore, each availability zone should use independent power, telecoms, and cooling infrastructure, which helps to contain the impact of any localised disruptions.

How to maximise the benefit

It's very easy to aim for more resilience than is needed, resulting in an unnecessarily expensive or complex design. Google's book on Site Reliability Engineering has a good discussion on not aiming too high with reliability. Correctly balancing load between multiple AZs is tricky and has some unexpected failure modes. Rather than reinventing the wheel, it's best to use services where the cloud provider already does this for you. For example, serverless components typically load balance across the AZs in a region automatically. Bulk data storage services also tend to have high resilience by default. Make sure you have a clear understanding of how your cloud provider provides resilience. What's their

equivalent of an AZ? How much will they do for you by default? Do any services have higher or lower resilience than the rest?

Outcome 13: Arbitrary capacity

A good cloud service should behave as if there is no hard limit on capacity. Whether it's the amount of compute, storage, or network capacity, the service should function as if capacity is limitless and can grow instantaneously. Removing capacity limits prevents several common problems, such as availability issues or outages while deploying security fixes. At the same time, cloud services have an inherent defence against capacity exhaustion denial of service attacks, since capacity is less likely to run out.

Having arbitrary capacity brings secondary benefits to other aspects of security. For example, having arbitrary storage capacity allows richer log and audit data to be retained for longer and to be stored with more context, creating opportunities for more complex analysis. This makes any investigations into abuse or attacks substantially more effective.

How to maximise the benefit

I've talked before about the security benefits of serverless. It can be much easier to gain the benefits of arbitrary capacity when using serverless components, as they tend to use resources more efficiently while responding to changes in demand more quickly. Furthermore, serverless compute tends to spread load over multiple AZs, whereas lower-level services like IaaS VMs tend to be bound to a specific AZ. This means that a capacity shortage in one AZ is more likely to affect you if you use IaaS than if you use serverless. Similarly, serverless storage services will use the vendor's storage capacity more effectively than running databases or filestores in IaaS.

Conclusion: addressing the big security challenges

This white paper has explained how using a good cloud service can bring significant security benefits to your organisation. Cloud services have been

designed with an array of security features in mind, and it's worth taking the time to find out what's available (and how to apply it your specific needs) in order to gain the most benefits. Good cloud services make advanced security accessible to **all** customers, not just those with large budgets.

The more you take the time to understand the cloud services available, the bigger these benefits will be. If instead you treat the cloud service as a generic provider of VMs, most of these benefits won't be realised. Some of these solve the **big security challenges** that we've struggled to address in pre-cloud environments (such as having an exhaustive inventory and a single, consistent approach to authentication, authorisation, logging, monitoring, and alerts). With these challenges effectively commoditised, you can focus your time and resources on solving business problems that are unique to your organisation, rather than constantly having to reinvent the wheel.

Finally, when selecting a cloud service, make sure that it meets your needs and helps you to secure your data. If you plan to implement or retain functionality that could be consumed as a service, make sure you have a clear understanding why you'd want to do this. Don't let anyone tell you the cloud is inherently insecure. If you choose a good cloud service - in line with our cloud security guidance - it's probably more secure than whatever it replaces.

Further reading

In addition to this white paper, readers maybe be interested in the following cloud-related publications, all of which are available on the NCSC website.

Cloud security guidance
Guidance on how to configure, deploy and use cloud services securely.

Software as a Service (SaaS) security guidance
Guidance for organisations looking to use, deploy, and understand the risks of adopting a range of popular Software as a Service (SaaS) applications.

Blog post: Why cloud first is not a security problem
Using the cloud securely should be your primary concern - not the underlying

security of the public cloud.

Blog post: Cloud backup options for mitigating the threat of ransomware
Why it's more important than ever to ensure your information is backed up securely.

PUBLISHED

13 November 2020

VERSION

1.0

WRITTEN FOR

Cyber security professionals

Large organisations

Public sector