Hybrid Clouds and Its Associated Risks

CSA cloud security alliance®

The permanent and official location for Hybrid Cloud Security Working Group is https://cloudsecurityalliance.org/research/working-groups/hybrid-cloud-security/

Hybrid Cloud Security Working Group (WG)

As businesses are developing rapidly, and Information Technology (IT) infrastructures constantly diversify, many cloud consumers find that a single public/private cloud or traditional on-premises datacenter is no longer able to meet service requirements in terms of costs, performance, scalability, security, resilience, regulations, and compatibility. Organizations are increasingly choosing hybrid cloud environments and services to meet their needs. Hybrid clouds take advantage of various clouds and traditional IT infrastructures and work systematically to benefit the users based on their service requirements. However, hybrid clouds pose different risks and thus bring on a different set of challenges to security. The WG aims to identify hybrid cloud security risks and countermeasures, helping users identify and reduce risk. Besides this, the WG also intends to provide suggestions on hybrid cloud governance, hybrid cloud threat profiles and hybrid cloud security evaluation, guiding both users and cloud service providers to choose and provide secure hybrid cloud solutions, and promote security planning and implementation.

Acknowledgments

Initiative Lead:

Zou Feng

Key Contributors:

John Barnes
David Chong
Christopher Hughes
Jean-Sébastien Mine
Michael Roza
Geng Tao
Saan Vandendriessche

Peer Reviewers:

Kevin Blackburn, FJ
Anjlica Dattatreya
Diego Diviani
Ankur Gargi
Jones Junior
Checri Loandos
Nya Murray
Adnan Rafique
Pierluigi Riti
Narudom Roongsiriwong

CSA Staff:

Jane Chow
Hing-Yan Lee
AnnMarie Ulskey (Design)
Haojie Zhuang

Table of Contents

1. Introduction

Cloud computing is flourishing. Hybrid clouds, especially, have been gaining more traction as cloud customers increasingly understand that using public clouds or private clouds alone poses certain limitations due to hardware or network restrictions. Hybrid clouds are now often the starting point for organizations in their cloud journey. International Data Corporations (IDC) IaaSView report in 2019 indicates that 52% of enterprises already have a hybrid cloud infrastructure in place[1]. Furthermore, Gartner predicts that by 2020, 90% of organizations will adopt hybrid cloud infrastructure management capabilities and services[2]. Multi-cloud convergence increases the complexity of risks in areas such as management, access control, data use, and service contracts. Security and compliance are among the issues that must be addressed in hybrid cloud use.

This paper aims to describe the concept and value of hybrid clouds, review its security risks, and highlight key use cases of hybrid clouds. A subsequent paper by CSA's Hybrid Cloud Security WG[3] will propose countermeasures to help users and cloud service providers (CSPs) mitigate and reduce security and compliance risks identified in this paper.

2. Hybrid Cloud Overview

2.1 Hybrid Cloud Concept

Server-side architecture typically falls into one of the following classifications: non-virtualized data centers, on-premises private cloud built by enterprises, hosted cloud by service providers for specific enterprises and public cloud by service providers for the public. There are inherent reasons such as cost, operation and maintenance (O&M) capabilities and regulatory requirements for different enterprises and applications to choose specific or different combinations of execution environments. Deploying hybrid clouds can congregate the best of different environments for users.

Under ISO/IEC 17788-2014[4], hybrid cloud is defined as a cloud deployment model that uses at least two different cloud deployment models (private, community, public; see Annex for definitions of each cloud deployment model). This definition will be adopted throughout this paper.

[1] https://www.idc.com/getdoc.jsp?containerId=prUS45625619

[2] https://www.gartner.com/en/newsroom/press-releases/2017-04-05-gartner-says-a-massive-shift-to-hybrid-infrastructure-services-is-underway

[3] https://cloudsecurityalliance.org/research/working-groups/hybrid-cloud-security-services/

[4] https://www.iso.org/obp/ui/#iso:std:iso-iec:17788:ed-1:v1:en

2.2 Business Value of the Hybrid Cloud for Enterprises

The hybrid cloud is an effective way to enjoy the benefits of public cloud without disrupting critical and core legacy services on private cloud.

Most enterprises' IT infrastructures contain legacy assets, including both software systems and hardware resources. Though dated, rigid, and oftentimes expensive to maintain, these legacy assets are often essential and core to the daily operations of the enterprise. As such, there is usually reluctance for the enterprise to disrupt or overhaul them. At the same time, enterprises want to maintain competitiveness and are compelled to continuously introduce new technologies and new cloud-oriented infrastructure architectures, especially front-end applications of non-core data systems which may also interact with core legacy services. Adapting and achieving this new infrastructure with legacy assets on their private cloud can prove unwieldy. Through a hybrid cloud solution, the enterprise can continue operating their private cloud and legacy services while tapping on the fast-evolving public cloud to keep up with competition.

A hybrid cloud may offer effective means to securely use cloud technologies.

In the process of deploying applications to the public cloud, enterprises could have security concerns that data on the public cloud may leak. The hybrid cloud is a solution that can help meet data protection requirements. Important data can be stored in a security-hardened zone, exposing only specific controllable interfaces to external applications in the public cloud. The hybrid cloud solution can also be used to deploy applications of different security levels/requirements to different clouds or infrastructure clusters, accessing them through external controllable API interfaces, thereby improving management efficiency.

A hybrid cloud can help enterprises leverage cloud resources at optimal costs.

If an unexpected peak usage occurs, server resources may become over-utilized and some services may have performance degradation, work improperly, or lack availability. In addition to simply adding hardware in the data center, there are also considerations of how to adjust the overall bandwidth in an end-to-end manner. Optimizing this is a time-consuming process. Using hybrid clouds is an efficient and cheaper solution that allows applications to scale transparently to end-users. There would also be savings in terms of security costs. By using hybrid clouds, increasing capacity or services related to non-sensitive data stored on a public cloud do not affect the cost of security related to sensitive data handled by private cloud infrastructure.

2.3 Hybrid Cloud Implementation

2.3.1 Layer 3 Network Interworking

To converge private and public clouds, the simplest way is to use a virtual private network (VPN) or dedicated private line to connect geographically separate networks. By interlinking the perimeter gateways of both virtual data centers, the enterprise can create a VPN connection between the virtual network in the public cloud and the network of the enterprise local data center.

This logical network consolidation provides seamless communication between sites that host virtualized services and enables applications to tap on greater resource pools. However, the disadvantage lies in that the resources of both ends may operate in differing physical cloud infrastructures, which then depend on the hybrid environment system owners for interoperability. Enterprises need to use different models to manage two independent clouds or utilize integration tools for orchestrating services across disparate cloud environments.

2.3.2 Multi-Cloud Management Enabled by Cloud Broker

This implementation is also known as a heterogeneous hybrid cloud, which means the architectures of parts of the hybrid cloud are different. A cloud broker provides a unified API and console interface, and adapts to different cloud APIs to implement unified resource management. It can be mounted to multiple clouds at the same time, but the type, quantity, and management supported by the cloud broker varies greatly. These tools include cloud management platforms (CMPs), cloud services brokers (CSBs), and other tools that provide an abstraction layer between clouds.

The upsides of using a cloud broker are increased cloud interoperability, cloud portability, and business continuity. However, it adds complexity to maintaining an organization's security requirements throughout the entire delivery chain.

2.3.3 Consistent Hybrid Cloud

Consistent hybrid cloud means the same architecture for the public cloud and private cloud. It can be considered as an extension of the public cloud in the local data center. At the IaaS layer, the private cloud and public cloud use the same/similar computing, storage, and network devices. The same architecture design enables seamless interconnection between applications and data. On the hybrid cloud management console, resource usage is uniformly displayed and managed in a unified manner. At the PaaS layer, the private cloud can synchronize PaaS capabilities in the public cloud, support more database functions and artificial intelligence (AI) capabilities, and support enterprises' application innovation based on core data. Because the PaaS platforms on both sides are the same, enterprises can implement offline development of applications, enjoy faster online deployment, and seamless interconnection. At the SaaS layer, enterprises can quickly develop and iterate their applications based on big data and AI, and provide more extensive ecosystem partner applications for enterprises.

This implementation enables services in the public cloud to be migrated to the private cloud (and vice versa) with few or no modifications. It uses unified development interfaces, unified O&M. and brings consistent cloud service experience to users. However, this implementation exerts higher requirements for CSPs. Currently, only large public cloud vendors offer this architecture, such as Alibaba Apsara Stack, AWS Outpost, Huawei Cloud Stack, and Microsoft Azure Stack.

3. Shared Responsibility in Hybrid Clouds

Cloud security is a shared responsibility between the CSP and the customer, where security responsibilities are distributed across the entire stack. As a general rule of thumb, security responsibility reflects the degree of control a party has over the architecture stack. Table 1 shows a typical shared responsibility model for either private or public cloud deployment models.

While the shared responsibility model in Table 1 typically holds true for pure private or public clouds, it is not as straightforward for the hybrid deployment model. In hybrid clouds, the responsibility distinction between customer and CSP tends to be blurred and fluid. This is attributed to the much more customized architecture of hybrid clouds, where it commonly varies between different organizations. As such, clear and precise contractual clauses should be drawn up for each customized hybrid cloud deployment to define the exact responsibilities of all parties involved.

	On-premises/ Private Cloud	Public Cloud		
		IaaS	PaaS	SaaS
Data	Customer	Customer	Customer	Customer
Applications	Customer	Customer	Customer	CSP
Runtime	Customer	Customer	CSP	CSP
Middle Ware	Customer	Customer	CSP	CSP
Operation System	Customer	Customer	CSP	CSP
Virtual Network	Customer	Customer	CSP	CSP
Hypervisor	Customer	CSP	CSP	CSP
Servers	Customer	CSP	CSP	CSP
Storage	Customer	CSP	CSP	CSP
Physical Network	Customer	CSP	CSP	CSP

Table 1

4. Risks, Threats and Vulnerabilities in Hybrid Clouds

Although the hybrid cloud environment seamlessly integrates private and public cloud, bringing onboard new IT capabilities to the hybrid cloud environment may introduce new security issues and concerns. Therefore, understanding and managing the associated risks must be a precondition before new capabilities are introduced to the cloud environment.

As a rule of thumb, the usual cloud security risks are generally applicable to hybrid clouds, for example, security threats listed in *Top Threats to Cloud Computing*[5] developed by CSA's Top Threats WG, risks listed in the report by ENISA[6]. In addition, for hybrid clouds, special attention needs to be paid to areas such as compliance and data security[7], which are of concern due to the interconnection between the public and private clouds. The following sections describe the common risks, threats and vulnerabilities that should be understood when adopting hybrid cloud.

4.1 Risks

4.1.1 Distributed Denial of Service Attack (DDoS)

The cost for attackers to launch a large capacity attack or application layer attack is considerably low compared to the severe impact it has in the cloud environment. The direct consequence of DDoS is network traffic congestion, which affects the quality of cloud services and ultimately the CSPs' reputation. In the hybrid cloud environment, DDoS attacks could presumably also disrupt "internal" communications between components of the hybrid solution.

4.1.2 Data Leakage

In the cloud environment, users' endpoints and clouds are connected through the open Internet. As such, users' data is at risk of leakage due to factors such as human errors or unauthorized access from man-in-the-middle attacks. For enterprises to continue to operate resiliently, critical data needs to be protected. This can be achieved through Data Loss Prevention (DLP) solutions, an important element of a broader security strategy surrounding data protection. When properly implemented, DLP can help address residual risks left behind by inbound security systems[8].

Also, when running hybrid clouds, service information may be (and typically are) exchanged between clouds via APIs. If left unprotected, malicious attacks may occur through the API interface in order to gain unauthorized access to data or modify configurations in the hybrid cloud, resulting in a breach of critical data or can cause a cloud/system outage.

[5] https://cloudsecurityalliance.org/research/working-groups/top-threats/

[6] https://resilience.enisa.europa.eu/cloud-security-and-resilience/publications/cloud-computing-benefits-risks-and-recommendations-for-information-security

[7] https://www-prev.pulsesecure.net/lp/hybrid-cloud-security-threats

[8] https://downloads.cloudsecurityalliance.org/initiatives/secaas/SecaaS_Cat_2_DLP_Implementation_Guidance.pdf

4.1.3 Perimeter Protection Risks

When constrained to only the private cloud, applications operate within clear network boundaries and security zones. However, once they migrate from private to the public cloud, original border security policies may be invalidated, affecting the applications' usage. This happens as corresponding security protection measures may not be on the same level for different clouds. In addition, the corresponding management responsibility boundary is now changed from single to multiple parties. Elastic scaling of applications could also introduce challenges to the division of security responsibility. There may also be gaps between the implementation of physical security across the different parts of the hybrid cloud (public vs private cloud/on-premises).

4.1.4 Compliance Risks

In hybrid clouds, achieving and maintaining consistent compliance is a huge challenge. Because data flows between on, and off-premises resources, it increases the difficulty to maintain and comply with governance frameworks in a hybrid model. For example, the global reach and storage capabilities of public clouds make it challenging to ensure that legal or regulatory requirements imposed on data, based on the country and region in which it resides, are complied with.

Compared to independent private or public clouds, hybrid clouds often comprise multiple service providers whose compliance capabilities are different. For sensitive operations such as cross-cloud data transfer, users would need to ensure close overall coordination and planning for the purpose of compliance management.

4.1.5 Misaligned Service Level Agreements (SLAs)

SLAs which define the service level objectives (SLOs) and service qualitative objectives (SQOs) are contracts between CSPs and customers. Generally, private clouds may have SLAs that are not as clear/stringent as that imposed when using public clouds. Furthermore, when dealing with multiple clouds, it is a challenge to align SLAs of different CSPs to deliver an overarching end-to-end service-oriented SLA to end-users. Different CSPs likely have their own tools, APIs, and SLAs that render it complex to link each underlying component for consistency. There could also be a delta in SLAs offered by public CSPs or private CSPs (co-location). In such situations, enterprises need to be aware of the differences so as to design & deploy their applications to fit the "right" stack.

4.1.6 Misalignment of Cloud Skill Sets

Public and private clouds may require separate skill sets, as they typically rely on different platforms to manage and monitor cloud services, which may not be directly applicable across clouds Additionally, each cloud may use specific management configurations and terminologies. For example, administration on AWS vs. Azure vs. private cloud can be vastly different. Considering that many security incidents on the cloud have been traced and attributed to insecure configurations or misconfigurations, the lack of appropriate cloud skill sets and knowledge can have dire consequences for enterprises trying to implement their cloud strategy. Security risks could arise if security configurations and controls are not applied correctly and consistently across different clouds.

Moreover, from a hybrid cloud system management perspective, enterprises need to coordinate different cloud skill sets and interconnect different platforms. For users, unified views and easy-to-use management tools are pertinent to manage hybrid clouds, simplifying management, and improving management efficiency. This requires proper cloud management platforms to centrally manage hybrid clouds and provide users with unified management, monitoring, and audit capabilities. Such needs are exacerbated if the enterprise uses multiple CSPs for different use cases.

4.1.7 Gap in Security Control Maturity

There could be misalignments or inconsistencies in the maturity of security controls in hybrid cloud setups. Oftentimes, public cloud environments are held to and have a higher level of security control maturity, or a more extensive security control catalog than typical private clouds. For example, some private cloud infrastructures may not be as conscientiously patched to the same levels as public clouds. This risk should be identified early and reviewed by senior management. If deemed unacceptable, the enterprise's security control catalog must be reviewed and updated where necessary, so that a consistent and standardized set of mature controls can be implemented. This will not only improve the security posture of the enterprise but could also result in financial cost optimization for the organization.

4.1.8 Comprehensiveness of Security Risk Assessment

Risk assessment can be challenging when evaluating hybrid cloud setups, where it is common for different providers to provide different parts of the infrastructure. For example, the private cloud may be owned and managed by the user organization. Risk assessment exercises may have been conducted separately for the private and public clouds rather than evaluated comprehensively as a whole. There could also be no detailed risk checks performed on the IT infrastructure and systems due a lack of management or technical tools that are applicable across the hybrid cloud. This affects the integrity of the assessment, causing security blind spots which may be misleading when assessing adherence or violation of security controls.

4.2 Threats

4.2.1 Malicious Insider

Not all internal users, be it the case for public or private clouds, are trustworthy. However, compared to internal threats in private clouds, malicious insiders in public clouds with weaker controls may be able to obtain cloud users' sensitive data, causing severe liabilities and disruptions to the enterprise. Also, in a hybrid set up, malicious employees or administrators may be able to compromise the public cloud using the private cloud as a conduit.

4.3 Vulnerabilities

4.3.1 Poor Encryption

While the clouds in a hybrid cloud architecture are individually subjected to regular data protection risks, the hybrid cloud as a whole faces higher risks due to the transit of data from one cloud environment to another. It is at the interconnection interfaces and pipes that data is most susceptible to theft or alteration if robust encryption is not employed. Users should ensure that state-of-the-art encryption is implemented between clouds to limit access even if the data is stolen.

4.3.2 Impacted Operational Processes

When an organization adopts a hybrid cloud architecture, processes have to be reviewed to evaluate if the current operational processes will continue to be applicable, or would be impacted or disrupted in any way -- especially when dedicated teams are formed to manage each cloud due to the different skill sets required. For example, in a hybrid environment, there may be limited management tools for provisioning or providing a single-pane-of-glass view for monitoring across the hybrid cloud.

4.3.3 Network Connectivity Breaks

Network connectivity between clouds in a hybrid cloud architecture is crucial for upholding SLAs, Business Continuity Plans (BCPs), and Disaster Recovery Plans (DRPs). Any slipups in establishing and maintaining this connectivity will drastically increase risks of service disruption, unavailability, and service quality degradation. There could also be single points of failures in the overall network architecture that may lead to widespread disruption of cloud services. For example, if backbone routing nodes lack redundancy, a single faulty backbone router is sufficient to cause an outage in the entire hybrid cloud.

4.3.4 Decentralized Identity & Credential Management

When the public and private clouds are integrated into a hybrid environment, using a centralized identity management solution is paramount to maintain a homogeneous identity lifecycle management process across multi-cloud environments. The lack of unified account management may cause account information inconsistencies between clouds, resulting in discontinuous log audits and failures to trace resource misuse. Identity authentication information must be shared in the hybrid cloud while retaining identity uniqueness, availability, and facilitating service applications.

4.3.5 Siloed Security Management

The private and public cloud components in a hybrid architecture may have disparate security management policies and processes. Independent deployment of respective clouds' policies may cause inconsistent management, resulting in management confusion and blind spots. Without unified management rules, the hybrid cloud can be easily compromised. Integrated programs with oversight over all components in the hybrid cloud should be adopted to address this issue.

5. Hybrid Cloud Use Cases

There are several common use cases of hybrid clouds[9&10]. All risks, threats, and vulnerabilities identified in the previous chapter are applicable to each of the use cases. However, for each of them, there are certain risks, threats, and/or vulnerabilities that are more prominent than the rest. It is also important to recognize that, as with other clouds, hybrid clouds also face the same threats as highlighted in CSA's Top Threats to Cloud Computing[11].

5.1 Workload Expansion (Bursting)

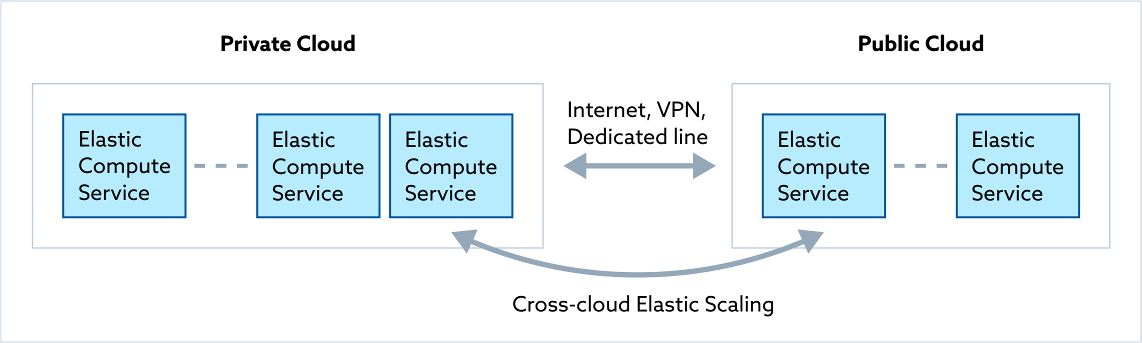

Figure 1

In this use case, applications are deployed in the private cloud. At any time when application access or usage spikes abruptly beyond the private cloud's capacity, enterprises will not be able to supplement hardware and resources to match demand in a timely fashion. With a hybrid cloud, the temporary capacity shortfall can be made up through workload expansion (bursting) from private to the public cloud. For example, large-scale e-commerce events (eg. Black Friday, Singles' Day) will see sharp spikes in the number of user visits, likely exceeding the capacity of any shopping platforms hosted on private clouds. Before this happens, the public cloud can be tapped on to compensate for insufficient computing resources, rather than expanding the private cloud and have the extra resources sit idle throughout lengthy non-peak periods.

Key risks, threats, and vulnerabilities pertaining to this use case:

- DDoS: Services suffer outages due to DDoS attacks.
- Data leakage: During inter-cloud data transmission, if the transmission processes/tools are not monitored or if no protection measures (eg.encryption) are taken, data leakage may occur.
- Impacted operational processes: In this use case, information processing of sensitive enterprise data may occur on the public cloud, which imposes stringent requirements for data security and privacy protection. A unified security management mechanism needs to be considered.
- Decentralized identity & access lifecycle management: Security control measures, such as identity authentication, authorization, and authentication management need to work seamlessly across the public and private clouds. A unified identification and verification mechanism needs to be considered.

[9] http://www.caict.ac.cn/kxyj/qwfb/bps/201906/t20190617_201254.htm
[10] http://www.caict.ac.cn/kxyj/qwfb/bps/201907/P020190704511581594525.pdf
[11] https://cloudsecurityalliance.org/research/working-groups/top-threats/

- Compliance risks: Cross-region or cross-country workload expansion increases compliance complexity in the hybrid cloud compared to compliance of a private cloud. Therefore, compliance management needs to be coordinated and planned in a unified and consistent manner.
- Network connectivity breaks: A network connection is the basis of data transmission. Breaks in network connectivity may result in the failure to burst, leading to service interruptions.

5.2 Backup

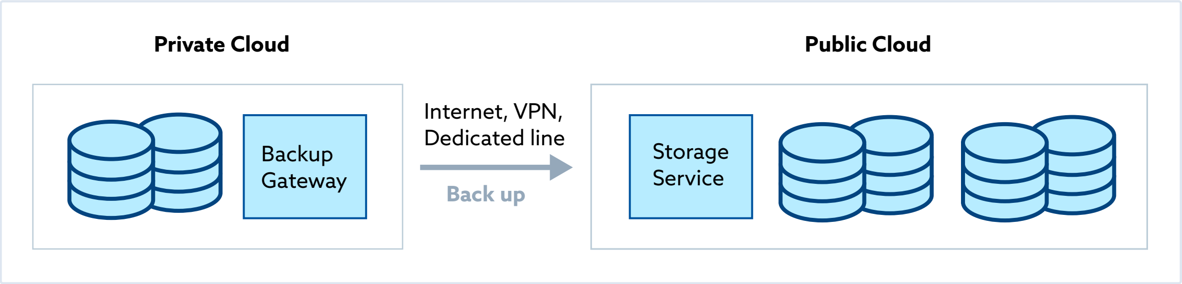

Figure 2

The purpose of backup is to redundantly store data or applications in secure and reliable locations. A common use case involves running applications on the public cloud, while the database containing sensitive data is stored in private cloud infrastructure for better assurance and data security. Data is then fetched by applications on the public cloud only when necessary. Such hybrid cloud setups are typically employed by applications that allow access to sensitive data such as healthcare records. This arrangement can also help meet compliance requirements of such industry sectors that need sensitive data to be stored in specific localities.

Key risks, threats, and vulnerabilities pertaining to this use case:

- Data leakage: During inter-cloud data transmission, if the transmission processes/tools are not monitored or if no protection measures (eg.encryption) are taken, data leakage may occur.
- Network connectivity breaks: A network connection is the basis of data transmission. Breaks in network connectivity may affect data integrity or even lead to failure of data backups that result in missed recovery time objective (RTO) / recovery point objective (RPO) targets.
- Compliance risks: Compliance capabilities of data backups may vary with CSPs. Closely monitor CSPs' overall compliance and management during data backups.

5.3 Disaster Recovery (DR)

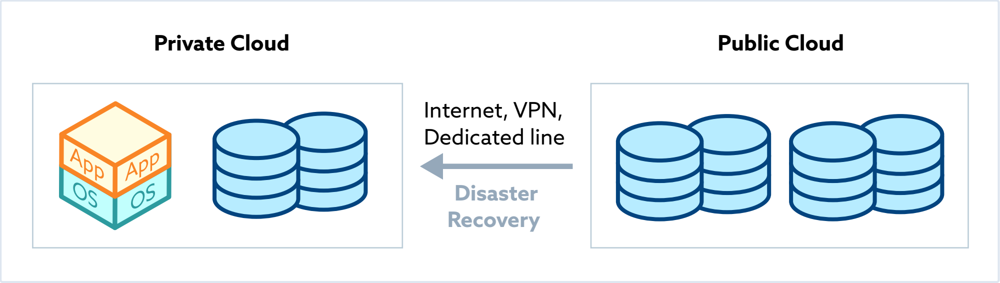

Figure 3

The construction of a data center that involves the implementation of physical security controls (eg. biometric locks, mantraps, HVAC, fire suppression systems) is a costly affair. Traditional two-site three-center DR can only be afforded by sizable enterprises such as large financial institutions. However, the hybrid cloud setup offers an economical solution that even small and medium-sized enterprises can easily achieve robust DR and backup. Data can be backed up to the public cloud, or active-standby/active-active DR can be implemented. With the right technical expertise and DR experience, coupled with O&M management of public CSPs, data can be restored quickly to ensure service continuity. In this use case, most enterprises tap on traditional cloud data center technology and solutions to implement dual-center DR and backup at low costs.

Key risks, threats, and vulnerabilities pertaining to this use case:

- Data leakage: During inter-cloud data transmission, if the transmission processes/tools are not monitored or if no protection measures (eg.encryption) are taken, data leakage may occur.
- Network connectivity breaks: A network connection is the basis of data transmission. Breaks in network connectivity may impact data integrity and lead to DR failure.
- Compliance risks: Compliance capabilities of CSPs may vary. Closely monitor CSPs' overall compliance and management during DR.
- Insufficient testing of DR plans & tools: Dry-runs are necessary to ensure that critical DR plans and components work during emergencies. DR tools are also often insufficiently tested. As such, DR may not achieve the expected results or even fail, causing irrecoverable loss.

5.4 Layered Deployment

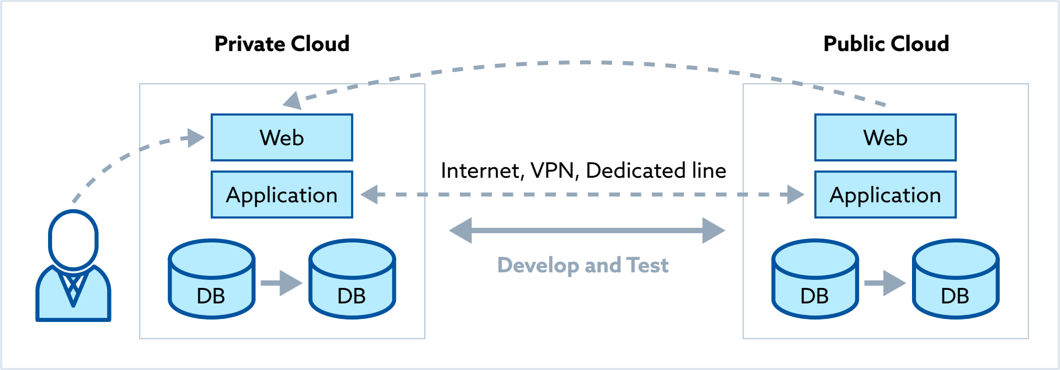

Figure 4

For enterprises with multiple regional headquarters/sites/branches, processing capabilities and access bandwidth can become bottlenecks if users are served centrally from the headquarters' data center. This is especially apparent during periods of demand spikes. A layered hybrid deployment will allow workloads to be shared and served from multiple and distributed infrastructures for better performance.

Key risks, threats, and vulnerabilities pertaining to this use case:

- DDOS: DDoS attacks can cause network traffic congestion within communication pipes connecting infrastructures of the branches and headquarters.

> - Perimeter protection risks: Because the application deployment environment is streamlined in this use case, security vulnerabilities or risks may propagate easily in the hybrid cloud. Care should be taken in border protection to prevent contagion (online attacks impacting the offline environment).
> - Network connectivity breaks: As the basis of data transmission, breaks in network connectivity may affect data integrity or even lead to service outages.

5.5 Application Container Technology

With the rapid development of container technologies, containers today can be easily migrated. In a hybrid cloud environment, application services encapsulated in containers can move from the private cloud environment to the public cloud and vice versa. In use cases such as automatic DR, traffic sharing, and elastic capacity expansion of cross-cloud applications, container technology helps to effectively reduce user access delay, improves application migration and capacity expansion efficiency, and enables applications to be deprovisioned quickly and flexibly.

Key risks, threats, and vulnerabilities pertaining to this use case:

> - Perimeter protection risks: Container migration between clouds may invalidate the original perimeter security protection measures, causing online (public cloud) attacks to spread to offline (private cloud) systems.
> - Gaps in cloud skill sets and security control maturity: In addition to security challenges of the container[12] itself, the hybrid cloud environment requires unified management of container resources and coordinated security policies. Security measures and monitoring must be implemented not only for containers but also for all application interfaces.

5.6 Extend New IT Capabilities

Public clouds, especially those provided by hyper-scale CSPs, offer a multitude of services and capabilities that can be leveraged to extend the capabilities and functionalities of existing on-premise systems. Examples are big data processing, AI /machine learning capabilities, serverless and API gateways among many others. Enterprises on a hybrid cloud can experiment and integrate required services to enhance their existing applications in a shorter timeframe, and with relatively less risk since the IaaS / PaaS capabilities have been proven prior. Compared to the traditional approach of having to invest significant resources to expand/rebuild existing IT systems and infrastructures, deploying a hybrid cloud environment puts existing on-premises resources to good use. This reduces costs, deployment time, and infrastructure complexity.

Key risks, threats, and vulnerabilities pertaining to this use case:

> - Compliance risks: Attention needs to be paid to safeguarding sensitive and critical data when processed in and across hybrid cloud environments.

[12] https://cloudsecurityalliance.org/artifacts/challenges-in-securing-application-containers-and-microservices/

- Impacted operational processes, misaligned SLAs, gaps in cloud skill sets and security control maturity: Unified security planning, SLAs, management tools, and processes are recommended to ensure secure interconnection between clouds.
- Non-unified APIs: This may cause inconsistent provisioning of cloud resources and security blind spots in management. Unified API design will help facilitate API maintenance and management.

6. Conclusion

The hybrid cloud is undoubtedly beneficial to the enterprise as it represents the best of both the public and private cloud worlds. But security and privacy in hybrid cloud environments remain a major concern, as they can be more challenging to manage than pure public or private clouds. This paper details hybrid cloud use cases and lists a number of risks, threats, and vulnerabilities typically faced in hybrid cloud deployments. Subsequent artifacts from the WG will focus on countermeasures and strategies to mitigate these weaknesses.

ANNEX - Definitions of Types of Cloud and Deployment Models

NIST Definitions:[13]

NIST **Private cloud**: The cloud infrastructure is provisioned for exclusive use by a single organization comprising multiple consumers (e.g., business units). It may be owned, managed, and operated by the organization, a third party, or some combination of them, and it may exist on or off-premises.

NIST **Community cloud**: The cloud infrastructure is provisioned for exclusive use by a specific community of consumers from organizations that have shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be owned, managed, and operated by one or more of the organizations in the community, a third party, or some combination of them, and it may exist on or off-premises.

NIST **Public Cloud**: The cloud infrastructure is provisioned for open use by the general public. It may be owned, managed, and operated by a business, academic, or government organization, or some combination of them. It exists on the premises of the cloud provider.

NIST **Hybrid Cloud**: The cloud infrastructure is a composition of two or more distinct cloud infrastructures (private, community, or public) that remain unique entities, but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load-balancing between clouds).

ISO Definitions:[14]

ISO **Private Cloud (3.2.32)**: deployment model (3.2.7) where cloud services (3.2.8) are used by a single cloud service customer (3.2.11) and resources are controlled by that cloud service customer (3.2.11).

ISO **Community Cloud (3.2.19)**: deployment model (3.2.7) where cloud services (3.2.8) exclusively support and are shared by a specific collection of cloud service customers (3.2.11) who have shared requirements and a relationship with one another, and where resources are controlled by at least one member of this collection.

ISO **Public Cloud (3.2.33)**: deployment model (3.2.7) where cloud services (3.2.8) are potentially available to any cloud service customer (3.2.11) and resources are controlled by the cloud service provider (3.2.15).

ISO **Hybrid Cloud (3.2.23)**: deployment model (3.2.7) using at least two different cloud deployment models (3.2.7).

[13] https://csrc.nist.gov/publications/detail/sp/800-145/final
[14] https://www.iso.org/standard/60544.html