



Best Practices Guide for Department of Defense Cloud Mission Owners

Version 1.0

Aug 2015

Last updated 2015-08-06

Developed by the

Defense Information Systems Agency (DISA)

For the Department of Defense (DoD)

IMPORTANT:

This Guide includes examples of specific products and vendor offerings. These examples are provided only as lessons learned and are NOT intended as recommendations. New versions of this Guide may include additional or different examples – no conclusions should be made about changes.

DISA does not endorse or recommend any specific product or vendor.

DISTRIBUTION A. Approved for public release: distribution unlimited.

“Cloud computing plays a critical role in the Department’s IT modernization efforts. Our key objective is to deliver a cost efficient, secure enough enterprise environment (the security driven by the data) that can readily adapt to the Department’s mission needs. The Cloud will support the Department’s JIE with a robust IT capability built on an integrated set of Cloud services provided by both commercial providers and DoD Components. We will use a hybrid approach to Cloud that takes advantage of all types of Cloud solutions to get the best combination of mission effectiveness and efficiency. This means in some cases we will use a purely commercial solution, which we have done with Amazon on public facing data, in others we will use a modified private Cloud hosted in commercial solutions, an example could be a shared federal or federal state government Cloud, and for our most protected data a DoD private Cloud that uses best industry practices.”

– Mr. Terry Halvorsen, DoD Chief Information Officer, statement to the House Armed Services Committee, Subcommittee On Emerging Threats & Capabilities (25 Feb 15)

Note:

This Best Practices Guide (BPG) is **NOT** DoD Policy, DISA Policy, a Security Requirements Guide (SRG), or a Security Technical Implementation Guide (STIG). It is a collection of Best Practices discovered during the DoD CIO Cloud Pilots effort for the benefit of the DoD Community.

The DoD Cloud Computing Security Requirements Guide is located at http://iase.disa.mil/Cloud_security/Pages/index.aspx. Compliance with the SRG is a requirement for all Cloud solutions, including commercial and government provided offerings.

Trademark Information:

Names, products, and services referenced within this document may be the trade names, trademarks, or service marks of their respective owners. References to commercial vendors and their products or services are provided strictly as a convenience to our users, and do not constitute or imply endorsement by DoD, DISA, or DISA Risk Management Executive of any non-Federal entity, event, product, service, or enterprise.

Table of Contents

1. Introduction	5
2. Initial Steps	6
3. Understanding the Shared Responsibility Model in Cloud Computing	9
4. Assessment and Authorization	10
5. IP Standards.....	11
6. Domain Name Service (DNS).....	13
7. Cloud Email	13
8. Storage Capacity, Partitions, and Backups	14
9. Achieving High Availability	18
10. Bastion Host.....	20
11. Useful Tips/Lessons Learned	21
Appendix A: Terminology.....	22

1. Introduction

Per the direction of the Department of Defense (DoD) Chief Information Officer (CIO), DoD examined the use of commercial cloud computing as a cost savings measure and for efficiencies in delivering highly available and scalable capabilities in support of the Warfighter. A 45-day study was commissioned to examine the balance between risks to the Department across the wide spectrum of computing needs and the costs of traditional security measures. The results of the study were twofold: 1) it produced the document called the “[DoD Cloud Way Forward](#)” and 2) it was the catalyst for the DoD CIO Cloud Pilots initiative. This Best Practices Guide (BPG) document is a collection of knowledge and experiences gained from the DoD CIO Cloud Pilots initiative, in particular DISA’s Information Assurance Support Environment (IASE) and U.S. Army’s DoD Environment, Safety and Occupational Health Network and Information Exchange (DENIX).

This BPG is targeted towards DoD Mission Owners who are planning to migrate an existing information system from a physical environment to a virtualized cloud environment.

Note: The Cloud Computing Security Requirements Guide (SRG) referenced throughout this BPG is located here:

http://iase.disa.mil/Cloud_security/Pages/index.aspx

Mission Owners must comply with the requirements of the SRG for any commercial cloud environment, including government and commercial hosted cloud services.

2. Initial Steps

Go to the cloud services support portal at the following link:

<http://disa.mil/Computing/Cloud-Services/Cloud-Support>

Click on the tab “How to Order.”

The screenshot shows the DISA Cloud Service Support portal. At the top, there is a search bar and social media icons. The navigation menu includes: About, Computing, Cybersecurity, Enterprise Services, Network Services, Mission Support, and Initiatives. The breadcrumb trail is: Computing > Cloud Services > Cloud Service Support. The main content area has a sidebar on the left with categories: COMPUTING, APPLICATION HOSTING AND VIRTUALIZATION, CLOUD SERVICES (with 'Cloud Service Support' selected), INFRASTRUCTURE, and MAINFRAME HOSTING. The main content area features a large image of a server rack with a 'LEARN MORE!' call to action. Below the image is the heading 'CLOUD SERVICE SUPPORT' and a row of tabs: Overview, Additional Information, How To Order, and Service Support. A paragraph below the tabs reads: 'DoD cloud consumers may seek support from the DISA Cloud Services Support office to fulfill their cloud service requirements. The DISA Cloud Services Support office uses the cloud service request form (CSR) to collect information from cloud consumers about their cloud service requirements and match the requirement with cloud provider capabilities.'

Click on the **Cloud service request form** highlighted in the text. This will navigate to the **Cloud Services Site**. At the top, there is a link to the Service Catalog. Mission Owner can use this link to check which Cloud Service Offerings (CSOs) have been granted Provisional Authorizations (PA). In addition, there is a link to the PA which contains specific information about the CSO. Next, the Mission Owner should evaluate the mission requirements versus the capabilities in the approved service offerings for goodness of fit.

Additionally, the Mission Owner must consider the desired scalability requirements. Mission Owners have more control over compute than bandwidth. Most Cloud Service Providers (CSPs) employ metered compute and bandwidth. Compute can be further divided into processor, memory, and storage.

- **Processor:** In a cloud based environment, the processing capability is provisioned across virtual CPUs. These vCPUs are hyper-threaded and are not mapped to a distinct physical core. In essence, the vCPUs are slotted on logical

processors and the physical core acts a scheduler of vCPUs' threads. It is important to note that increasing the number of vCPUs does not always equate to an upsurge in performance since the scheduler has to coordinate all the actions of the vCPUs. Therefore, if a workload requires 8 vCPUs, it may be more suitable to break the workload across two 4 vCPUs VMs or four 2 vCPUs VMs via a load balancer.

- **Memory:** Virtual Memory is measured in GBs. Monitoring usage is easy and if the utilization approaches 90% of the available memory for 80% of the time, it is time to increase the memory.
- **Storage:** This by far is the least expensive component of the Cloud. To answer the question "How much storage do I need?" please refer to the [Capacity](#) section below.

Understanding the Impact Levels of the data, as detailed in the DoD Cloud Computing SRG is critical. It is highly recommended that Mission Owners review and understand the SRG, as well as the associated security controls attributed to each Impact Level.

After researching the different types of Cloud options and determining the best CSO, it is important to discuss the effort with the AO thoroughly.

Last, for easy reference Table 1 below provides an Impact Level comparison regarding key security control requirements and Table 2 provides a notional division of security risk between CSP and Mission Owner based on type of service being acquired.

IMPACT LEVEL	INFORMATION SENSITIVITY	SECURITY CONTROLS	LOCATION	OFF-PREMISES CONNECTIVITY	SEPARATION	PERSONNEL REQUIREMENTS
2	PUBLIC or Non-Critical Mission Information	FedRAMP v2 Moderate	US / US outlying areas or DoD on-premises	Internet	Virtual / Logical PUBLIC COMMUNITY	National Agency Check and Inquiries (NACI)
4	CUI or Non-CUI Critical Mission Information Non-National Security Systems	Level 2 + CUI-Specific Tailored Set	US / US outlying areas or DoD on-premises	NIPRNet via CAP	Virtual / Logical PUBLIC COMMUNITY Strong Virtual Separation Between Tenant Systems & Information	ADP-1 Single Scope Background Investigation (SSBI)
5	Higher Sensitivity CUI Mission Critical Information National Security Systems	Level 4 + NSS & CUI-Specific Tailored Set	US / US outlying areas or DoD on-premises	NIPRNet via CAP	Virtual / Logical FEDERAL GOV. COMMUNITY Dedicated Multi-Tenant Infrastructure Physically Separate from Non-Federal Systems Strong Virtual Separation Between Tenant Systems & Information	ADP-2 National Agency Check with Law and Credit (NACLC) Non-Disclosure Agreement (NDA)
6	Classified SECRET National Security Systems	Level 5 + Classified Overlay	US / US outlying areas or DoD on-premises CLEARED / CLASSIFIED FACILITIES	SIPRNET via CAP	Virtual / Logical FEDERAL GOV. COMMUNITY Dedicated Multi-Tenant Infrastructure Physically Separate from Non-Federal and Unclassified Systems Strong Virtual Separation Between Tenant Systems & Information	Favorably Adjudicated SSBI SECRET Clearance NDA

Table 1, Impact Level Comparison (Source: DoD Cloud Computing SRG)




Table 2, Notional Division of Security Inheritance and Risk (Source: DoD Cloud Computing SRG)

3. Understanding the Shared Responsibility Model in Cloud Computing

In migrating workloads to the Cloud, the CSP will provide compute, bandwidth, and storage capabilities. It is important to note, that the CSP will provide security mechanisms, but the Mission Owner is responsible for activating, deploying, and managing the ones listed in Table 2 as *you manage*. The Mission Owner is also responsible for including security requirements in the Request for Proposals, contract, and relevant Task Orders.

For commercial Infrastructure as a Service (IaaS) offerings, for example, the Mission Owner is responsible for the security of the following components.

- Virtual Machine
- Operating systems
- Applications
- Data in transit
- Data at rest
- Databases
- Credentials to include Private Keys
- Adhering to DoD Policies and configurations i.e. STIGs
- Vulnerability Compliance Reporting

For all workloads, as a result of recent OMB directive M-15-13¹ (released 8 June 2015), the directive “requires all publicly accessible Federal websites and web services only provide service through a secure connection.” Last, the Mission Owner is required to encrypt the data in transit utilizing FIPS 140-2 compliant encryption IAW with the Application Security and Development STIG (V-6136).

TIP: Create a table/spreadsheet that lists the responsible organizations/ individuals and their specific responsibilities.

¹ <https://www.whitehouse.gov/sites/default/files/omb/memoranda/2015/m-15-13.pdf>

4. Assessment and Authorization

The Mission Owner **MUST** STIG the servers in the Cloud virtual environment.

The Mission Owner **MUST** deploy Host Based Security to the virtual machines unless the capabilities are provided by the CSP as part of a Software as a Service (SaaS) offering.

The Mission Owner **MUST** patch the virtual machines to include the Information Assurance Vulnerability Management (IAVM) bulletins.

The Mission Owner **MUST** get an Interim Authority to Operate (IATO) or Authority to Operate (ATO) for mission information systems hosted in the Cloud. The Provisional Authorization is specific to the CSO and is leveraged by the mission Authorizing Official (AO) as part of the risk decision process.

For the most part, the same process that is used to accredit a system that is hosted in a DoD Data Center is followed when accrediting the virtual environment. DoD has moved from the DoD Information Assurance Certification and Accreditation Process (DIACAP) to DoD Risk Management Framework (RMF)². It is recommended that the Mission Owner coordinate with the Information System Security Manager (ISSM) and seek guidance from his or her Authorizing Official (AO) in determining what is needed to assess and accredit the system.

The following are typical artifacts that should be captured and built as part of standing up the environment:

- Network Architecture Diagram showing the virtual servers, the virtual enclaves, the subnets, IPs, and how they are connected
- Software Listing
- Site Security Plan approved by the ISSM
- Security Content Automation Protocol (SCAP) Benchmark Scores
- Vulnerability Compliance Report

² http://www.dtic.mil/whs/directives/corres/pdf/851001_2014.pdf

5. IP Standards

To determine the appropriate IP requirements, it is helpful to be familiar with the Classless Inter-domain Routing (CIDR) Block system. The CIDR is a set of IP standards that are used to create unique identifiers for networks and individual devices. The IP addresses allow particular information packets to be routed to specific devices such as VMs. An abbreviated CIDR block table³ is provided below.

IP/CIDR	Δ to last IP addr	Mask	Hosts (*)	Class
a.b.c.d/32	+0.0.0.0	255.255.255.255	1	1/256 C
a.b.c.d/31	+0.0.0.1	255.255.255.254	2	1/128 C
a.b.c.d/30	+0.0.0.3	255.255.255.252	4	1/64 C
a.b.c.d/29	+0.0.0.7	255.255.255.248	8	1/32 C
a.b.c.d/28	+0.0.0.15	255.255.255.240	16	1/16 C
a.b.c.d/27	+0.0.0.31	255.255.255.224	32	1/8 C
a.b.c.d/26	+0.0.0.63	255.255.255.192	64	1/4 C
a.b.c.d/25	+0.0.0.127	255.255.255.128	128	1/2 C
a.b.c.0/24	+0.0.0.255	255.255.255.000	256	1 C
a.b.c.0/23	+0.0.1.255	255.255.254.000	512	2 C
a.b.c.0/22	+0.0.3.255	255.255.252.000	1,024	4 C
...				
a.b.0.0/16	+0.0.255.255	255.255.000.000	65,536	256 C = 1 B
...				
a.0.0.0/8	+0.255.255.255	255.000.000.000	16,777,216	256 B = 1 A

* For routed subnets bigger than /31 or /32, two reserved addresses need to be subtracted from the number of available host addresses: the largest address, which is used as the broadcast address, and the smallest address, which is used to identify the network itself. In addition, any border router of a subnet typically uses a dedicated address.

Table 3, CIDR Block Table

For, Impact Levels 4 and 5, the Mission Owner will need to request DoD IP space from the DoD Network Information Center (NIC): <https://www.nic.mil>. The example on the next page pertains to IaaS Impact Levels 4 and 5. Impact Level 2 environments will use IPs from the CSO's IP space.

³ <http://whatismyipaddress.com/cidr>

Let's walk through an example to determine how many IPs the Mission Owner might need. For this example, let's assume a Mission Owner is setting up 7 servers in the Cloud.

1. Primary Active Directory
2. Secondary Active Directory
3. Web Server A
4. Web Server B
5. Database Server
6. File Server
7. Bastion Host (more on this in next section)

Mission Owner might assume they need 7 IPs. But, that would not be enough. Let's look for compliance with the Web Server STIG, in particular placing the Web Server in a separate subnet of other assets (V-2242). The Mission Owner will want to create at least two Subnets labeled as public and private. Each CSO will have minimum CIDR block limits for creating subnets.

For example, AWS Virtual Private Cloud (VPC) has a minimum CIDR block limit of /28⁴ and the block limit for Microsoft Azure Virtual Network (VN) is /29⁵. So if the Mission Owner used AWS VPC, the block limit would be /28 which is 16 IPs. For two subnets, that is 32 IPs. Remember that for each subnet, two IPs are reserved for broadcast and identity, resulting in 28 IPs. In a second example with Azure VN, the block limit is /29, which is 8 IPs. For two subnets, that is 16 IPs. Again, remember that for each subnet, two IPs are reserved for broadcast and identity and Azure reserves two IPs for its services per subnet, resulting in 8 IPs.

It is important to note to that the Mission Owner must look beyond just the VMs in determining how many IPs are needed. Since the Mission Owner is creating the enclave in IaaS offering and as such, establishing the subnets, he or she must also think **about growth**. Last, the Mission Owner should consider the IPs required in deploying security devices such as an Assured Compliance Assessment Solution (ACAS)⁶ Scanner. If so, additional IPs may be required.

TIP: When constructing a virtual enclave, please take the time to properly label all items, in particular the subnets that are private or public. This will help others to know what was done, as well as assist in creating a network diagram.

⁴ http://docs.aws.amazon.com/AWSEC2/latest/APIReference/API_CreateSubnet.html

⁵ <https://msdn.microsoft.com/en-us/library/azure/jj156074.aspx>

⁶ <http://www.disa.mil/Cybersecurity/Network-Defense/ACAS>

6. Domain Name Service (DNS)

Once a Cloud Mission Owner acquires IP space, it is time to plan for DNS. The Cloud IPs may not be in the same address space as the Mission Owner's DNS provider. Therefore, the Mission Owner needs to discuss with the CSP how to obtain and manage both the "A" records (forward lookup) and "PTR" (reverse pointer or reverse DNS) records as needed.

The "A" record (forward lookup) maps a domain name to an IP address. The PTR does the opposite by mapping the IP address back to the domain name. PTR records are used to confirm the server at the originating IP is legitimate and approved. Most importantly, **a PTR record is required to send email from your cloud application to the internet.**

The A record should be provided upon request from the DNS provider.

TIP: Recommend coordination with your DNS provider in the planning phase of the cloud project

7. Email from Cloud Mission Space

If the application (from its own E-mail Service) sends NIPRNet-sourced email to the internet, the Mission Owner needs to configure the cloud application to use the Enterprise Email Security Gateway (EEMSG). DISA publishes a guide called EEMSG Phase II Migration Guide. The document is FOUO. Please contact your DISA Enterprise Information Services Customer Management Executive (CME) representative to get a copy. A list of CMEs is here

<http://disa.mil/About/~-/media/Files/DISA/Services/Computing/CMEContact.pdf>

Keep in mind that a DNS PTR record is required to use EEMSG service.

8. Storage Capacity, Partitions, and Backups

Cloud computing, when done right, can provide substantial increase in storage agility. Provided below are three disk space recommendations for those new to virtual servers.

1. *Deployment of a Microsoft Windows Server 2012 R2 into a virtual machine.* According to TechNet⁷, the minimum space required is 32 GB. However this does not take into consideration any additional roles or features you would have your server take on. Additionally, it is a best practice to **have two separate partitions** one for the Operating System (OS) and one for the Data. In a starter set up, recommend 40 - 50 GB for the OS space (assuring that the virtual memory paging file has enough room to breathe) and 50 GB for the Data. If this is a web server, you may want to bump up the Data to 75 - 100 GB. Please note that to be in compliance with the Internet Information Services (IIS) 7.0 STIG rule (V-3333) place the wwwroot (or other web directories) on the Data partition:

“The web document (home) directory is accessed by multiple anonymous users when the web server is in production. By locating the web document (home) directory on the same partition as the web server system file the risk for unauthorized access to these protected files is increased. Additionally, having the web document (home) directory path on the same drive as the system folders also increases the potential for a drive space exhaustion attack.”

2. *Deployment of a SQL Server into a virtual machine.* According to MSDN⁸, the minimum space required is 6 GB, which is relatively small. It is recommended to have the OS have 50 GB of space (assuring that the virtual memory paging file has enough room to breathe) and a separate partition with 50 GB of space for Data. Please note in accordance with the Microsoft SQL Server 2012 Database Instance STIG rule (V-40951):

*“Install SQL Server software using directories separate from the OS and other application software library directories.
Relocate any directories or reinstall other application software that currently shares the DBMS software library directory to separate directories.”*

One last approach includes three separate partitions: one for the OS, one for the Data, and one for the Audit Logs. This simplifies locking down the auditing logs, in compliance for protecting audit information from any type of unauthorized access in accordance with the Microsoft SQL Server 2012 Database Instance

⁷ <https://technet.microsoft.com/en-us/library/dn303418.aspx>

⁸ <https://msdn.microsoft.com/en-us/library/ms143506.aspx>

Security Technical Implementation Guide (V-41016). This approach does not require a separate partition for audit logs, but imagine the scenario in which your audit logs grow and you need more disk space. The Mission Owner could easily swap out that partition without affecting the OS or Data. This would benefit a web server as well.

3. *Deployment of a LINUX OS into a virtual machine.* There are many varieties of LINUX. For this example, let's consider Red Hat Linux and SUSE Linux Enterprise Server. First, Red Hat Linux⁹, the minimum storage requirement is 4 GB. This is pretty small. It is recommended to have 10 – 15 GB instead. Please remember that the minimum specifications on the Red Hat site do not take into consideration the auditing requirements as described in the Red Hat Enterprise Linux 6 STIG. The Red Hat Enterprise Linux 6 STIG requires the following partitions:

Partition	Rule
/home	V-38473
/tmp	V-38455
/var	V-38456
/var/log	V-38463
/var/log/audit	V-38467
/var/www (Optional, Web Servers, also recommend changing webroot to non-default path)	From Web Server STIG V-3333
/opt (Optional, Commercial Applications)	Author Recommendation

Note: The above does not mean the system needs 7 distinct volumes. Multiple partitions can exist on one volume.

Table 4, Red Hat Linux Partitions

⁹ https://access.redhat.com/documentation/en-US/Red_Hat_Directory_Server/8.2/html/Installation_Guide/Installation_Guide-RHEL-Requirements.html

Secondly, SUSE Linux Enterprise Server¹⁰, the recommendation is 2 GB available disk space (8.5 GB for all patterns). Again, this is relatively small. One common reason for choosing to deploy SUSE Enterprise Server is because Apache 2.2, MySQL 5.5, PHP 5.3, and Ruby 1.8.7 are installed. It is recommended the OS have 15 – 20 GB of storage and the Data have 10 GB. The Mission Owner should utilize YaST to configure the VM to have two partitions. Then configure Apache 2.2 to store the data files on the separate mounted Data partition. **Remember to properly apply permissions to whatever you create IAW the STIG.**

Now let's switch to another related topic – **backups**, which is a Mission Owner's responsibility using IaaS. In the planning stages, the Mission Owner should think about the frequency of backing up the data. There are multiple approaches available to achieve this. One method is to create weekly images of all of the VMs. This may be viewed as overkill. Another approach is to take snapshots of specific volumes. If the Mission Owner follows the process of separating OS from Data, the OS snapshot may be monthly or right after a patch cycle completes and your Data would be more frequent. To illustrate the importance of deliberate backup,

Rackspace in their Best Practices for Cloud Backup¹¹ defines a backup strategy according to three factors: criticality, size, and data churn. Several points are worth summarizing:

- **Criticality** is the most important factor to consider when making backup decisions. The more critical the file is to business operations, the more often one should back this file up¹¹.
- **Size** is an important consideration if the speed of backups/restores is important. Large files take longer to backup and to restore. It may be wise to backup large files less frequently¹¹.
- **Data churn** is the last variable to consider, and perhaps the trickiest to handle. Files that change often invalidate blocks that have been stored previously. Depending on criticality, it may still be wise to snapshot files with high data churn and backup those snapshots¹¹.

¹⁰ <https://www.suse.com/products/server/technical-information/>

¹¹ http://www.rackspace.com/knowledge_center/article/best-practices-for-Cloud-backup

Back up less often files that have...	Back up more often files that have...
Lower Criticality	Higher Criticality
High Data Churn	Low Data Churn
Larger Size	Smaller Size

TIP: “Do not make decisions about backup frequency lightly. If you try to backup or restore files more frequently than the backup engine can keep up with, anomalous behavior may result.”¹¹

There is no wrong approach. Please remember to document the process in the Site Security Plan which is required by Application Security and Development STIG.

9. Achieving High Availability

Regardless of the CSO, there are specific mechanisms to achieve high availability in the Cloud. The Mission Owner must design for a robust architecture.

Before going further, it is important to understand that the Risk Management Framework (RMF) takes into consideration three security objectives¹² (FIPS 199) for information and information systems: **Confidentiality (C)**, **Integrity (I)**, and **Availability (A)**.

Under this method, each of the three objectives is rated **High (H)**, **Moderate (M)** or **Low (L)** for each information system. An approach will differ if a system is rated as [C:M; I:M; A:H] versus rated as [C:L; I:L; A:L]. The takeaway is to work with your ISSM and AO office in determining the information system's rating.

Let's consider five different failures that may affect the availability of a specific system/application in the Cloud.

1. *Application Failure*: To illustrate, SharePoint Server is running and the web server service fails. One solution is to build in fault tolerance into the application and define behaviors for application or system errors. This includes information that must be logged in the error logs. Additionally, do not forget test and evaluation of the Cloud information systems.
2. *Server Failure* (including misconfiguration): One solution involves spreading the workload across multiple servers and utilizing a **load balancer**. AWS calls this Elastic Load Balancing¹³ and Azure calls it Traffic Manager¹⁴. One issue is the frequency of data replication. The Mission Owner must work with their stakeholders in determining the appropriate level of data replication.
3. *Data Center failure* (including zone failure): Most CSOs offer **multiple zones** [data in another data center]. One solution is to split the load into different zones. AWS calls this Availability Zones¹⁵, Azure calls these Regions¹⁶, other CSOs may have different names but the concept is all the same.


¹² <http://csrc.nist.gov/publications/fips/fips199/FIPS-PUB-199-final.pdf>

¹³ <http://aws.amazon.com/elasticloadbalancing/>


¹⁴ <https://azure.microsoft.com/en-us/documentation/articles/traffic-manager-overview/>

¹⁵ <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.html>

¹⁶ <http://azure.microsoft.com/en-us/regions/>



Example of AWS Availability Architecture¹⁷



Example of Azure Region¹⁸

4. *Whole Cloud Failure.* It may be that a system is mission critical enough, for example a RMF Rating of [C:M; I:M; A:H]. In this case, the Mission Owner could place backup information into another CSO as described in section 5.12 of the Cloud Computing SRG.
5. *Network Failure.* There will be times of network congestion within the users' enclave. There will also be Internet Congestion. Mission Owners should develop appropriate plans of actions for such events.

TIP: Regardless of what Availability techniques are used in the Cloud, the Mission Owner must not forget to establish some sort of notification mechanism. Some CSOs offer built-in HTTP Status Code 200¹⁹ checks against a file on the web server. There may be other notification mechanisms that check the health of the VM itself.

¹⁷ <http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/using-regions-availability-zones.html>

¹⁸ <http://blogs.msdn.com/b/hanuk/archive/2014/12/19/building-mission-critical-systems-using-Cloud-platform-services.aspx>


¹⁹ <http://www.w3.org/Protocols/rfc2616/rfc2616-sec10.html>

10. Bastion Host

When implementing a Cloud environment, Mission Owner should utilize as many “Defense in Depth” approaches as possible. One of the least expensive security systems is a Bastion Host. From Committee on National Security Systems (CNSS) Instruction No. 4009²⁰, a Bastion Host is “a special purpose computer on a network specifically designed and configured to withstand attacks.” The computer generally hosts a single application, for example a proxy server, and all other services are removed or limited to reduce the threat to the computer.

A cloud-based Bastion Host sole purpose is to provide **administrative access to other VMs**. It is recommended that a box with 1 – 2 vCPUs and 4 GB of memory is available. When not in use the Bastion Host should be turned off. Additionally, if using Active Directory and establishing a domain, recommend that the Bastion Host is NOT a member of the domain. Therefore, if by chance the Bastion Host is compromised, the attacker will not have direct access to the domain.

Lastly, do not forget to **STIG the Bastion Host** and **keep it patched**.



²⁰ www.ncsc.gov/publications/policy/docs/CNSSI_4009.pdf

11. Useful Tips/Lessons Learned

There are several things that are helpful in knowing before launching the Cloud environment.

1. Regardless of what CSO is used for IaaS, the CSP will offer a plethora of choices for instance types. Do not get caught up in worrying about selecting the wrong type. It is easy to change the instance type in the Cloud. In most cases, the Mission Owner will get two choices: (1) the ability to dynamically decrease or increase the instance size on the fly or (2) the ability to create an image of the instance to redeploy it in a different size. In standing up the cloud environment, the Mission Owner will likely change instance sizes multiple times. Furthermore, if the Mission Owner has a tight budget and has to choose between vCPU and Memory, go with Memory.
2. When deploying a web front end server, it can be discovered by web crawlers like Google and Yahoo! Upon discovery, the web crawlers can hit the site hard and eventually bring the site down. The remedy is the placement at the root site level of a robots.txt file, which deters or deflects web crawlers. For more information on how to build a robots.txt including syntax and format, please go here: <http://tools.seobook.com/robots-txt/>
3. Estimating bandwidth usage-based billing can be difficult. There may be little basis for estimating the requisite bandwidth. One project estimated 500 GB in traffic per month, but soon realized it was closer to 2 - 5 TB per month. To estimate for bandwidth utilization, it is recommended a Mission Owner multiplies his or her initial estimate by 4. Since bandwidth is metered, if a Mission Owner significantly overestimates the usage, they will not be paying for it. Recommend that the Mission Owner reviews bandwidth usage at least quarterly.

DISA welcomes Mission Owners and Industry Partners to submit lessons learned found during their Cloud implementations. Please send a note of interest to the disa.meade.re.mbx.disa-commercial-cloud@mail.mil.

Appendix A: Terminology

This Guide uses the following terminology:

Cloud Service Provider (CSP): A company that offers cloud computing service components, specifically Infrastructure as a Service (IaaS), Software as a Service (SaaS) or Platform as a Service (PaaS) offerings.

Cloud Service Offering (CSO): A specific capability or set of capabilities offered for consumption. These offers must have a FedRAMP approval and have a DoD Provisional Authorization (PA).

The Federal Risk and Authorization Management Program, (FedRAMP): A government-wide program that provides a standardized approach to security assessment, authorization, and continuous monitoring for Cloud products and services. For more information go here: <https://www.fedramp.gov/>

Infrastructure as a Service (IaaS)²¹: The capability provided to the consumer to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying Cloud infrastructure but has control over operating systems, storage, and deployed applications; and possibly limited control of select networking components (e.g., host firewalls).

Platform as a Service (PaaS)²¹: The capability provided to the consumer to deploy onto the Cloud infrastructure consumer-created or acquired applications created using programming languages, libraries, services, and tools supported by the CSP. The consumer does not manage or control the underlying Cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly configuration settings for the application-hosting environment.

Software as a Service (SaaS)²¹: The capability provided to the consumer to use the CSP's applications running on a Cloud infrastructure. The applications are accessible from various client devices through either a thin client interface, such as a web browser (e.g., web-based email), or a program interface. The consumer does not manage or control the underlying Cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.

²¹ NIST Definition of Cloud Computing, Special Publication 800-145.

Community Cloud²²: Provisioned for exclusive use by a specific community of Mission Owners from organizations that have shared concerns (e.g. mission, security, 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.

Private Cloud²²: Operated solely for a single Mission Owner, whether managed internally or by a third-party and hosted on or off premises. Private Clouds are usually virtualized Cloud data centers inside a Mission Owner's firewall, or they may be private space dedicated to an organization within a CSP's data center. In private Cloud computing, access to the Cloud is limited to internal users. Additionally, in private Cloud computing, users may still have to purchase the internal hardware and software, implement and manage the solution and thus do not benefit from lower upfront capital costs and less hands-on management. Of course, these costs can still be mitigated somewhat by creating more efficient environments through virtualization – or by taking the off premises route to private Clouds by asking the CSP to provide dedicated resources for critical data or applications.

Public Cloud²²: The delivery of Cloud services (e.g., software applications) over the Internet by a third-party provider to the general public. It exists on the premises of the CSP and usually includes virtualization for more efficient deployment of shared resources.

Hybrid Cloud²²: Any combination of external public CSOs and internal resources to create a solution. Hybrid cloud computing implies a significant integration or coordination between the internal and external cloud environments. It may exist on or off premises. In the completely off-premises hybrid model, a CSP can supply a shared virtual environment along with private, dedicated, data storage space all on the CSP's premises.

²² Best Practices for Negotiating Cloud-Based Software Contracts, <http://www.esi.mil/download.aspx?id=4783>