AWS Cloud

Services Description

Version 1.1

TeraGo

Contents

TeraGo Offered AWS Services 3

Amazon Simple Storage Service (S3) 3
Amazon Elastic Block Store (EBS) 3
Amazon Elastic File System (EFS) 4
Amazon Glacier 4
Amazon Elastic Compute Cloud (EC2) 5
Amazon Relational Database Service (RDS) 5
Amazon Aurora 6
AWS Application Discovery Services (ADS) 6
AWS Database Migration Services (DMS) 7
AWS Server Migration Services (SMS) 7
Amazon Virtual Private Cloud (VPC) 7
AWS Direct Connect 8
Amazon Elastic Load Balancing (ELB) 8
Amazon CloudWatch 9
AWS Systems Manager 9

TeraGo Managed Services 10

Managed Backup and Recovery 10
Disaster Recovery as a Service (DRaaS) 11
Additional DR Test 12
Managed High Availability Service 13
Managed OS Patching 14
Managed Core Monitoring 14
Managed Network 15
Standard Monthly Reporting 16
Advanced Monthly Reporting 16
Billing Optimization Reporting 17
Managed Firewall Service 18
Managed DDoS Protection (Advanced) 18

TeraGo Professional Services 19

Cloud Assessment 19
Cloud Migration Services 20

TeraGo

TeraGo Offered AWS Services

Amazon Simple Storage Service (S3)

Amazon S3 is a highly available, durable, and scalable storage service that is ideal for storing an unlimited number of any type of objects (documents, video, HTML/CSS/JS, etc.). Amazon S3 is designed to be scalable, which means as your business needs change and your data grows, Amazon S3 will automatically expand the storage capacity to easily meet your business needs. Amazon S3 is ideal for storing backup files and offers 99.999999999% (eleven nines) of durability, which means if you store 10,000 objects with Amazon S3, you can on average expect to incur a loss of a single object once every 10,000,000 years.

With Amazon S3 you can store multiple versions of a single object and have full access to easily and quickly restoring any version of an object at any time. Each object in S3 is stored inside resources called buckets and each object can be 5 terabytes in size. S3 buckets enable you to store unlimited number of objects. You can easily control access to the buckets and grant appropriate permissions to users in terms of who can create access or delete objects. By default, only the owners of buckets and objects have access to the data. With Amazon S3 you can upload and download data via SSL endpoints using HTTPS protocol. Additionally, you can view access logs at the bucket and object level, providing you with additional control over your underlying data.

Amazon S3 comes in two tiers - Standard and Standard-Infrequent Access, which are designed to help you balance your business and performance needs with cost of the storage. Standard tier is recommended for frequently accessible data while the Standard-Infrequent Access, as the name implies, is a cost-effective option designed for storing data where access to the data is infrequent in nature.

You can define data lifecycle management policies to automatically move data from the Standard tier class to Standard-Infrequent access to Amazon Glacier. With this automation in place, you can further improve cost-savings associated with your storage.

Amazon Elastic Block Store (EBS)

Amazon EBS is a highly available, performant, and low-latency persistent block storage attached to the running Amazon Elastic Compute Cloud (EC2) instance in the same Availability Zone. Amazon EBS volumes are persisted independently from the life of an EC2 instance, which means if you terminate or shut down your EC2 instance, the EBS volume will be available to be attached to another instance. Amazon EBS volumes offer high performance for applications that require quick and continuous access to the stored data and can easily be scaled up or down based on your business needs. An EBS volume is automatically replicated within its Availability Zone to provide you with high availability and durability while protecting your data against component failure.

Amazon EBS provides encryption of data as it moves between EC2 instances and EBS storage as well as data encryption at rest using the Amazon-managed keys or keys created by you using the AWS Key Management Service (KMS).

Amazon EBS volumes can be up to 16 terabytes in size and each volume can only be attached to a single running EC2 instance while there could be multiple EBS volumes attached to a single EC2 instance. With Amazon EBS you can create snapshots of the volumes and store them in S3 for high resiliency; it also supports point-in-time snapshots of modified blocks.

Here are few options available in terms of the EBS volume types:

- **Provisioned IOPS (io1)** – offers high performance and low latency making it suitable for transactional workloads
- **General Purpose SSD (gp2)** – good for wide variety of transactional workloads but the focus is to help you balance price and performance
- **Throughput Optimized HDD (st1)** – low cost storage option that is ideal for frequent data access and throughput intensive workloads
- **Cold HDD (sc1)** – ideal for less frequently accessed workloads and offers lowest price for the storage

Amazon Elastic File System (EFS)

Amazon EFS is a highly available, durable, and scalable file system to be used with Amazon Elastic Compute Cloud (EC2) instances. It is designed to provide consistent low-latency access and enables you to stage your backups to S3 for reduced cost. Unlike the EBS volumes, a single Amazon EFS volume can be mounted to multiple EC2 instances, providing the flexibility of a common data source across multiple applications. By leveraging AWS Direct Connect service, you can access EFS file systems from your on-premises servers.

Amazon Glacier

Amazon Glacier is a low cost and long-term storage option for storing infrequently accessed data, making it ideal for backup and archival storage. With Amazon Glacier there is no need to keep a tape backup in your data center or offsite; there is also no need to provision storage capacity that ultimately doesn't get used. Amazon Glacier provides a reliable and cost-effective way to store your long-term data that is seldom accessed.

By default, only the owner has access to the data stored within Amazon Glacier; however, you can leverage AWS Identity and Access Management (IAM) to grant specific users access to the stored data. The data is encrypted on the server side using the 256-bit (AES-256) encryption.

There are a few options available for data retrieval from Amazon Glacier:

- Expedited with expected retrieval time between 1 to 5 minutes
- Standard with expected retrieval time between 3 to 5 hours
- Bulk with expected retrieval time between 5 to 12 hours

Amazon Elastic Compute Cloud (EC2)

Amazon Elastic Compute Cloud (EC2) provides secure, resizable compute capacity in the cloud. EC2 provides a wide selection of instance types optimized to fit different use cases. Instance types comprise varying combinations of CPU, memory, storage, and networking capacity and give you the flexibility to choose the appropriate mix of resources for your applications. Each instance type includes one or more instance sizes, allowing you to scale your resources to the requirements of your target workload.

On these instances, you can choose from different operating systems such as Windows or Linux (Red Hat, Ubuntu, Amazon Linux, etc.) and pre-installed software packages such as SQL Server (Standard, Web, Enterprise) to which you will be given root access and the ability to interact with them as you would any machine.

When selecting your instance, there are various other options available to select from, for example, whether you want to reserve compute capacity on a 1 to 3-year term and benefit from the reduced monthly cost while paying a portion of the term cost upfront. Reserved instances are a great way to save money, if you have a steady-state usage scenario of the duration of the term. If you anticipate a varying usage pattern for your EC2 instances, then it will be better to choose On-Demand instances instead to ensure optimal performance at all times. In any event it’s important to understand whether this option makes more sense for your business vs the On-Demand instances with no upfront fee or term-contract but a higher monthly cost over reserved instances.

Amazon Relational Database Service (RDS)

Amazon Relational Database Service (RDS) provides a cost-effective way of running relational database instances in the AWS cloud. There are a number of database instance types available to choose from, for example, memory optimized instance, performance optimized, or provisioned I/O. Depending on your application needs we can help you select the appropriate database instance type across various database engines, including Amazon Aurora, PostgreSQL, MySQL, Microsoft SQL server, Oracle and MariaDB.

TeraGo provides migration services to help you move your on-premises databases to the cloud to help you reduce cost while enhancing application availability and reliability for your mission-critical production workloads. Amazon RDS offers a high availability architecture with built-in automated fail-over.

By default, Amazon RDS provides automated backups which enable point-in-time recovery of your database instance. It will backup your database and transaction logs and store them for the retention period you define. Besides the automated backups, you can create snapshots of your database and store them in Amazon Simple Storage Service (S3). These snapshots are kept until explicitly delete.

With Amazon RDS you can provision multi-AZ database instances which will synchronously replicate your data to a secondary instance in a different availability zone.

You can encrypt your Amazon RDS databases using Amazon Key Management Service (KMS). The data stored at rest, along with backups, read replicas, snapshots can be encrypted. Any data in transit is secured via SSL.

Amazon Aurora

Amazon Aurora is a high performance relational database service available in the AWS Cloud and it is the fastest growing service in AWS history. It is a MySQL and PostgreSQL compatible relational database that combines the performance and availability of high-end commercial databases with the simplicity and cost-effectiveness of open source databases.

Aurora is up to five times faster than standard MySQL databases and three times faster than standard PostgreSQL databases.

All Amazon Aurora database instances must be created within an Amazon Virtual Private Cloud (VPC). Amazon Aurora uses SSL (AES-256) encryption to ensure the connection from your application to the database is secure. You have the option to encrypt the database using Amazon Key Management Service (KMS). Amazon Aurora utilizes Amazon's encryption capabilities. This also applies to the automated backups, snapshots, and replicas in the same cluster.

AWS Application Discovery Services (ADS)

AWS Application Discovery Services (ADS) are used as part of the migration effort to capture information about the on-premises data centers. As part of the AWS migration planning effort, ADS helps you capture system inventory, performance data and dependencies. ADS agents are deployed on the source hosts and provides support for both Windows and Linux. The data captured is securely sent and stored in the Application Discovery Services (ADS) store. The information saved in AWS cloud can be extracted as CSV for further analysis.

AWS Database Migration Services (DMS)

AWS Database Migration Services (DMS) helps you migrate your databases from your on-premises data centers to the AWS cloud. You can migrate to and from commercial and open-source databases. With DMS, your applications will continue to work while replication is taking place in parallel. DMS supports both types of migration, such as homogenous (from and to same database engine) and heterogeneous (from and to different database engines).

AWS Server Migration Services (SMS)

AWS Server Migration Services (SMS) is an agentless service which helps you migrate your on-premises workloads to AWS Cloud. AWS SMS enables you to create and manage replication schedules and provides you with full visibility into migration progress of each workload. With AWS SMS you can replicate your VMWare virtual machines to AWS by saving them as Amazon Machine Images (AMI). These AMIs can then be launched as EC2 instances in the AWS cloud.

With AWS SMS you can migrate Windows Server 2003, 2008, 2012, and 2016 as well as Windows 7, 8, and 10. AWS SMS also supports migration of Red Hat Enterprise Linux (RHEL), SUSE/SLES, CentOS, Ubuntu, Oracle Linux, Fedora, and Debian Linux operating systems.

Amazon Virtual Private Cloud (VPC)

With Amazon Virtual Private Cloud you launch AWS resources in a virtual network that is private and logically isolated section of the AWS cloud. You have full control over your virtual networking environment, including selection of your IP address range, creation of subnets, and configuration of route tables and network gateways. Amazon VPCs can span over multiple Availability Zones in a region. They support both IPv4 and IPv6 for secure and easy access to your resource and applications. If you have a security conformance requirement to logically isolate your workloads in the cloud, you can setup / operate multiple VPCs and set up VPC peering to configure connections between the individual VPCs.

Within the Amazon VPC you can setup public and private subnets, for example you can add your web server in a public subnet with a public / Elastic IP, while leaving the database and the app server in a private subnet. This will ensure the resources running inside the private network are secure and not reachable from the internet. Enabling communication between the resources running within public and private subnets or across VPCs is supported and easy to configure via Network Access Lists. The same applies with enabling the instances running inside a private subnet to reach the internet to download patches; this can be setup by configuring a Network Access Translation (NAT) Gateway inside the public subnet, which will behave as a proxy for the resources in the private network that are allowed to communicate to the internet.

Additionally, you can establish a private network between AWS and your corporate office. This can be accomplished by setting up Direct Connect service and the appropriate gateway endpoints on each side. By doing so, any communication between AWS and your corporate office is completely private and does not touch the internet. You also get the benefit of much higher bandwidth.

AWS Direct Connect

With AWS Direct Connect, you can establish a private connection from your on-premises data center to your Amazon Virtual Private Cloud (VPC). With a private connection, you can reduce costs, increase bandwidth, and provide more consistent network performance vs an internet-based connection. Now you can transmit data to AWS securely over private network vs sending it over the internet.

AWS Direct Connect integrates with Amazon Elastic Compute Cloud (EC2), Amazon Virtual Private Cloud (VPC), Amazon Simple Storage Service (S3), and other AWS native services.

Amazon Elastic Load Balancing (ELB)

Amazon Elastic Load Balancing (ELB) helps direct traffic (HTTP / HTTPS / TCP) to multiple targets, such as Amazon EC2 instances, containers, and IP addresses. Amazon ELB can be configured to utilize instances within a single Availability Zone or across multiple Availability Zones. Amazon ELB continually monitors instances to ensure it will only direct traffic to ones that are healthy.

There are three types of load balancers available – Application Load Balancers, Network Load Balancers, and Classic Load Balancers. Load Balancers are integrated with Auto Scaling and automatically scale up or down based on the incoming traffic.

AWS Auto Scaling

With Auto Scaling, you can ensure your application is always available and as the demand changes Auto Scaling will automatically scale up or down the Amazon EC2 capacity. This ensures an optimal use of the compute capacity while also ensuring you only pay for what you use instead of over paying for unused capacity or suffering due to performance degradation due to under provisioned capacity.

With Auto Scaling, you can also define conditions to ensure you never go above or below the capacity limits of your fleet. It will automatically increase the number of running EC2 instances during high demand and reduce the instance count during low demand. By doing so, it will help you contain costs while ensuring your application is always available to meet demand.

Auto Scaling will automatically detect unhealthy instances and replace them with new instances should the situation arise. If configured for multi-zone deployment Auto Scaling will automatically launch new instances in other availability zones within the region and balance the instances across zones.

AWS Auto Scaling is enabled by Cloud Watch and does not carry any additional pricing.

Amazon ELB and AWS Auto Scaling are essential services to achieve high availability within your AWS cloud.

Amazon CloudWatch

Amazon CloudWatch provides you with visibility into resource utilization and helps you define alarm thresholds, send notifications, and respond to events. For example, you can trigger Auto Scaling to spin up new EC2 instances when your CPU utilization exceeds 80% or spin down EC2 instances when CPU utilization drops below 20%. With CloudWatch you can do things like log aggregation, monitoring, and troubleshooting.

AWS Systems Manager

AWS Systems Manager helps you manage your EC2 and on-premises instances. It helps automate day-to-day tasks so you can focus on the business strategy vs spending time updating and maintaining your systems. It provides full visibility and control of your infrastructure on AWS and enable you to efficiently group AWS resources and take appropriate actions collectively. For example, you can group resources, such as Amazon EC2 instances, Amazon RDS, or Amazon S3 buckets, which simplifies management of applications.

TeraGo

TeraGo Managed Services

Managed Backup and Recovery

From AWS to AWS

Each workload that you run has important data that is used to serve your customers. This data is critical to the operations of your business and data loss can cause a lot of distress and negatively affect your business operations and your customers. Loss of data can occur for many reasons, e.g. underlying hardware failure, natural disaster, human error, etc. Protecting your most valuable asset should be the first business priority.

TeraGo provides the Managed Backup and Recovery service in AWS for efficiently backing up your AWS resources. This service includes the following backup targets:

- EC2 instances (including EBS volumes)
- EBS volumes, regardless of being attached to EC2 instances
- RDS databases
- RDS Aurora clusters and Redshift Clusters

Snapshots will be taken at the block level which will ensure only the modified disk blocks are backed up. Our experienced staff will work with you to define a backup schedule and retention period with the schedule frequency ranging from minutes to months to better meet your business needs. This approach will ensure great flexibility in Recovery Point Objective (RPO). We can also help you define data lifecycle management policies to ensure optimal use of the AWS storage space.

The public cloud backup service is powered by N2WS and includes following services.

- Backup AWS resources and store back into AWS
- Policy-based backups for backup targets and snapshots
- Management and hosting of backup storage to support this service
- Proactive support and alerting on missed backups or other issues
- Backup will be stored in Amazon S3 to minimize costs
- Troubleshoot any failed backups
- Monthly reporting on backup data
- 24/7 Service Desk to support client reported incidents

Customer Responsibilities:

- Test and validate backups to ensure the backed up data is complete and valid upon restoration
- Define data backup and retention policies
- Submit ticket to TeraGo Service Desk for any change request

Not included in this service:

- Configure data lifecycle policies to move backed up data to low cost storage, i.e. S3 to Glacier – this service is available as a onetime Professional Services engagement.

Disaster Recovery as a Service (DRaaS)

From On-premise VMware/Hyper-V, TeraGo Cloud, or AWS

To support your critical business functions against any disruption caused by unplanned events, such as natural disasters, ransomware attacks or server crashes, a good IT resiliency strategy is essential and should be part of the overall business continuity plan. To reduce significant impact to business operations, revenue, customer experience and productivity, it becomes crucial to continue to have all systems operational during a disaster or any other downtime.

For businesses whose virtual (VMware/Hyper-V) IT infrastructure resides *on-premise*, or within *TeraGo's cloud*, or on *AWS*, TeraGo provides IT resiliency / Disaster Recovery for the following use-cases:

Production site	***Recovery site options***
(Customer) On-Premise	TeraGo's Private Data Center
	AWS
TeraGo's Private Data Center	TeraGo's Private Data Center
	AWS
AWS	TeraGo's Private Data Center
	AWS

Fully managed services as part of DRaaS, enable failover as well as failback of applications and data within pre-set recovery time objectives (RTO) and recovery point objectives (RPO) to allow business continuance. Virtualized workloads within VMware, or Microsoft Hyper-V environments are supported.

This managed service includes:

1. Solution Design – Determination and documentation of all solution design aspects including infrastructure, network & storage requirements and key replication/recovery software utilized within the DR solution. Customer will be provided with a detailed runbook, documenting the following:
   - System and network (LAN, WAN) configuration
   - Application details and interdependencies
   - Authorization and access details
   - Roles and responsibilities for TeraGo managed services & customer IT
   - Failover steps – activation of the DR plan
   - Workload failback steps
   - TeraGo support information including escalation steps
   - Recovery Time and Recovery Point Objectives (RPO/RTO)

2. Implementation - Deployment of disaster recovery solution as determined within the 'Solution Design' step above, as well as provisioning of necessary system administration and operational support for:
   - Disaster recovery software & underlying virtualization technologies

- Storage, network and firewall configurations

Initial failover & failback testing is to be conducted within 30 days of solution implementation. For every subsequent year, one additional failover & failback test is included at no extra charge. Note that the DR test will run within a period of 30 days or less.

3. Monitoring & Recovery - Recovery of customer workload upon disaster declaration. This includes:
   - Monitoring of customer workloads for scheduled replication tasks
   - Ensuring replication activities meet recovery point objectives
   - Failing-over customer workloads (as described in the runbook) in case of disaster declaration and within predefined RTOs
   - Completing failback of customer workloads once primary location is available

4. Maintenance & Testing - includes:
   - Disaster Recovery test in conjunction with customer's IT staff. Customer is to notify TeraGo 90 days in advance of desired DR test date. As described previously, one instance of a managed DR test is included per calendar year at no extra charge.
   - Technical support and troubleshooting for all disaster recovery related software and hardware components, including firewalls, DRaaS software and hypervisor layer.
   - Quarterly maintenance of DR runbook capturing any changes or updates required due to software, infrastructure or personnel changes.

Customer responsibilities:

- Declaration of an outage scenario. Confirmation that workloads are required to be brought online at secondary (DR) site.
- Management, monitoring & technical support for software applications installed *within* guest VMs.
- Advising TeraGo of updates or changes to systems/applications/configurations, etc., via ticketing system.

Additional DR Test

Additional DR tests can be purchased per virtual machine (VM) and are executed on a per availability group (i.e. group of VMs). DR tests are performed according to the most current runbooks as maintained by TeraGo. Execution reports will be provided by TeraGo and will outline the actual RTOs and RPO, with potential remedies proposed.

Managed High Availability Service

High Availability (HA) is critically important if you need a reliable and fault-tolerant system in the cloud that delivers optimal performance and guaranteed uptime. Especially if your business operates around the clock, you want to ensure your customer experience is not negatively impacted in way. Therefore, a high availability architecture becomes a fundamental requirement to deliver optimal customer experiences and maximizing revenue potential of your organization. We will ensure your production site continues to function and minimizes any risk of downtime due to infrastructure failure or variable spikes in the demand. Our Disaster Recovery as a Service (DRaaS) services will ensure protection against unexpected disasters that could take down your entire site and significantly affect your business operations, application data, and financials.

AWS services are highly reliable and cost-effective options for achieving optimal performance and uptime for any type of workloads operating in the AWS cloud. As part of the infrastructure setup, TeraGo will deliver a scalable and flexible architecture that will scale up or scale down based on the demand. Our Solution Architects are highly trained and certified on the AWS services and will ensure your systems are properly architected in the cloud.

Depending on your use cases, our high availability setup may include following AWS services:

- **Elastic Load Balancing (ELB)** – network and/or application load balancers, depending on the business needs. Automatically route traffic to healthy EC2 instances and replace unhealthy instances running behind the load balancers while the site continues to be functional
- **Auto Scaling** – to automatically scale up or down based on the demand
- **Elastic IP** – keep same IP address when new instances are launched or existing instances are replaced within an AWS Region. These IP addresses are tied to the AWS account and not to the actual EC2 instance
- **Elastic Block Store (EBS)** – offers persistent off-instance storage volumes delivering high durability tied to a single Availability Zone (AZ). Point-in-time snapshots can be created and stored in Amazon Simple Storage Service (S3) for cross-AZ replication within a region. Since new EBS volumes can be launched from a snapshot, this would serve as a backup in case any disk failures, host-level issues, or any issues affecting Availability Zones (AZ)
- **AWS Availability Zones (AZ)** – geographically isolated locations in an AWS Region. Launching instances as a multi-zone deployment will ensure protection against any failure in a single Availability Zone (AZ)

Included in this service:

- 24x7x365 Environment monitoring and management
- Initial Environment creation
- Security group and policy configuration

Customer Responsibilities:

- Management of customer data
- Advising of updates or change requests via TeraGo's ticketing system

Managed OS Patching

Patch management is an optional service whereby server management tasks (up to and including the Operating System) are conducted by the service provider. This approach makes customer resources available to focus on their core competencies and their business. Customers receive tailored automated alerting, dashboards and reporting features.

Included in this service:

- Recommendation of operating system updates and configuration modification with Customer concurrence to apply update
- Minor upgrades to the server operating system, which includes service packs, minor version upgrades
- Operating system patch updates, including security and integrity patches as required and agreed to by the Customer

Customer responsibilities:

- Application management, troubleshooting and testing, including customer provided software
- Database and information management and troubleshooting
- Schedule patch maintenance with TeraGo support
- Advising of updates or change requests via TeraGo's ticketing system

Managed Core Monitoring

With the Managed Core Monitoring service, TeraGo will ensure your systems are healthy by constantly monitoring for anomalies using Amazon CloudWatch. We will detect and resolve potential issues that may impact your system performance. Through this Managed Service, our goal is to minimize any type of disruption to your business. For this reason, performance and availability of your applications and AWS infrastructure will be monitored around the clock by TeraGo. When issues are detected, our system will automatically trigger email notifications and alert our NOC team as well as the customer. Appropriate members of the team will be assigned to immediately investigate the identified issue.

Included in this service:

- 24x7x365 environment monitoring and management
- Collect and track metrics
- Proactive alarm investigation and take appropriate actions to resolve potential issues
- Dashboards and reporting via the Customer Service Center (CSC)
- Set alarms and take appropriate actions to resolve potential issues
- Service Provider retains and manages root access
- Management of problem determination and resolution of server management activities
- Management of existing operating system configuration, by modifying configuration files, documenting system configuration, and controlling access to system configuration files

- Management of operating system files, by creating, maintaining and deleting volumes and directory structures, modifying file system sizes, verifying mount point availability, repairing defective file systems, and modifying file system permissions
- Management of operating system processes (e.g., continuously running system subtasks); by refreshing processes as required, establishing start up sequences, and changing process priorities as appropriate
- Maintain tools for server management to enable installs, modifications and removals
- Automatic ticket creation and proactive electronic customer notification for detected hardware and resource limit alerts
- Initiate critical “call-out” notification of up to five (5) Customer supplied contacts for Severity 1 issues (as defined by a service being "hard" down or a critical impact to a customer's business operation with no possible workarounds for the customer, its users, or the service provider);
- Evaluation of planned changes to the server environment and advise Customer of any requirements to support such changes

Customized monitoring dashboard and historical reporting accessible through portal.

Included in this service:

- Application aware monitoring
- In-guest OS monitoring
- Log aggregation and analysis

Customer responsibilities:

- Configuring and managing operating system
- OS patching unless subscribed to Managed OS patch management
- Installed software applications
- Advising of updates or change requests via TeraGo’s ticketing system

Managed Network

With the Managed Network service, TeraGo will setup the network service and ensure the network meets all customer expectations and SLAs. At the core of this service, TeraGo will monitor:

- Network availability / uptime – availability of systems at all time
- Network utilization – address any hardware failure
- Managed DirectConnect – Utilizing TeraGo’s Canada wide network backbone to minimize latency and enhance performance
- VPN Configuration – Configure site-to-site VPN’s on AWS and provide customer configurations for various hardware options

Standard Monthly Reporting

After your resources are provisioned, it becomes crucial to monitor the usage of these resources to ensure all resources are operational and running optimally. These reports provide insights into:

- CPU utilization
- Memory utilization
- Number and type of storage, and the amount utilized in GB
- Network utilization

The usage data will be provided in a real-time reporting dashboards in the TeraGo Customer Service Center (CSC).

Advanced Monthly Reporting

In addition to the Standard Monthly Reporting, you can leverage the Advanced Monthly Reporting service which includes best practices and recommendations pertaining to your environment in AWS.

A TeraGo specialist will gather the appropriate information and organize it in a readable format. These advanced reports will be delivered to you in an email on monthly cadence. A follow up meeting will be scheduled with you to discuss the report content and document any resulting actions.

- **EC2 Instance Right Sizing Report** – provides EC2 instance type recommendations based on the last 30 days utilization of memory, CPU and network.

- **Best Practices Reports** – providing high-priority best practices on Security, Cost, Availability, Usage, and the AWS Trusted Advisor.

- **Billing & Cost Management Report** – this report shows spend summary by service based on a blended cost type, e.g. EC2, KMS, S3, etc. The report displays spend for last month, month-to-date, and forecast until the end of the month.

Included in this service:

- Monthly reports as stated above
- Walkthrough of the reports on monthly cadence
- 24/7 Service Desk to support client reported incidents

Customer Responsibilities:

- Contact the assigned TeraGo specialist and request details of a particular report or AW Service under the framework of the reports included in this service

Billing Optimization Reporting

With the Cost Optimization Reporting, TeraGo will ensure you are provided with insights into how your environment in AWS is being used over a period of time. As the demand changes you want to make sure your environment is right sized so you can maximize the benefits. This information will provide you with an opportunity to continuously optimize your environment from cost perspective without compromising on performance and availability of your AWS infrastructure services.

TeraGo specialist will gather the appropriate information related to cost optimization and organize it in a format to simplify the decision making process. The Cost Optimization Reports will be delivered to you in an email on monthly cadence. A follow up meeting will be scheduled with you to discuss the report content and document any resulting actions.

- **Cost Savings Report** – this report includes information on the following and the possible monthly savings:
    - Idle resources that are adding up to the cost
    - Unused resources
    - Reserved purchase recommendations

- **EC2 RI Purchase Recommendations (by instance)** – this report details which RI option to purchase on a 1 or 3 year term and offering class as well as the region. The report will detail total saving based on No Upfront, Partial Upfront, or All Upfront payment schemes.

- **Historical Monthly Billing Summary** – this report will provide a month to month cost information across all services that are being consumed. Customer can request details of a particular service during the monthly touch point with the TeraGo specialist.

- **Inventory Summary Report** – this report includes information on the AWS services being consumed across the account and the amount of each service that is being consumed. Additionally it will include the AWS bill (month-to-date). Customer can request details of a particular service during the monthly touch point with the TeraGo specialist.

Included in this service:

- Monthly reports as stated above
- Walkthrough of the reports on monthly cadence
- 24/7 Service Desk to support client reported incidents

Customer Responsibilities:

- Contact the assigned TeraGo specialist and request details of a particular report or AW Service under the framework of the reports included in this service

Managed Firewall Service

Your data is the most valuable asset that you have and safeguarding it from unauthorized visitors and potential malicious threats should be one your top business priorities.

Securing your systems and continuously monitoring and maintaining them is resource intensive and requires expertise to effectively deploy necessary measures to safeguard against unauthorized access and prevent costly breaches.

TeraGo provides advanced firewall protection along with 24x7 monitoring and responding to security alerts to protect your IT assets and networks. We minimize the business risk associated with security breaches. The goal of managed firewall service is to continuously monitor the network traffic to identify and resolve potential threats before any harm can be done. You focus on your business while we focus on protecting your systems from external threats.

Our managed firewall services may include the following AWS products and services:

- Web Application Firewall (WAF) to prevent unauthorized access to the application resources. Note: there is additional charge for using WAF.
- Network Address Translation (NAT) Gateway to prevent external access to backend servers
- Configuring subnets and security groups
- Security Group management and changes – up to 10 changes a month. Additional changes will incur additional charge
- Patching and updating
- 24/7 monitoring and issue resolution

We have trained security professionals to monitor and manage your enter system including networks, application and services. Put your IT staff to better use by leveraging TeraGo as your Managed Services provider.

Customer responsibilities:

- Notify TeraGo of application security, ports and network requirements

Managed DDoS Protection (Advanced)

Requires a minimum of 1 year term commitment

As part of our advanced DDoS protection service, TeraGo will take necessary measures to protect your applications from various types of DDoS attacks that are targeted to take down your systems and/or steal your data. Our goal is to ensure your systems are running and operating normally while the potential threats are taken care of systematically and seamlessly through the deployment of effective tools and architecture best practices.

As part of the AWS cloud setup, we will ensure Web Application Firewall (WAF) are deployed along with AWS Shield Advanced to give you the maximum protection. AWS WAF is designed to block or allow web requests based on header inspection, IP addresses, etc. Furthermore, AWS WAF will provide traceability into the request being made so that the firewall rules can be updated to ensure latest and greatest protection of your AWS resources.

This service provides coverage for the following threats:

- DDoS to protect against HTTP floods that are intended to take down your system, for example:
    - Adding unnecessary stress on systems such as firewalls, load balancers via state-exhaustion DDoS attacks
    - Volumetric DDoS attacks – congest networks by directing a large volume of network traffic
    - Application-layer DDoS attacks – malicious requests designed to consume application resources, for example, Get Requests, DNS query floods, etc.)
- Other application specific attacks such as SQL injection and application exploits
- Bad bots, website crawlers, content scrapers, etc. that are designed to steal your data

With AWS Shield Advanced, we will ensure monitoring and detection is always on and provide advanced L3/L4 & L7 DDoS protection, attack notification and reporting. AWS Shield can sit with the Application Load Balancer, Classic Load Balancer, Amazon CloudFront, and Amazon Route 53.

Our service includes:

- Fully managed service
- Enhanced DDoS detection and monitoring for Elastic Load Balancing (ELB), CloudFront (content distributions), and resources attached to an Elastic IP, e.g. EC2 instances
- Includes Web Application Firewall (WAF) when protected with AWS Shield Advanced
- PCI-DSS certified
- Experienced staff
- 24/7 support

TeraGo Professional Services

Cloud Assessment

There are many challenges that come up when running on-premises workloads. These challenges can range from being mild with little impact to the business or severe that cost millions of dollars. If you are running on-premises workloads then you may have experienced some or all of the following challenges:

- Difficulty of upgrading and maintaining on-premises workloads
- Capacity planning – cost of under / over utilized resources
- Lack of or insufficient visibility into resource utilization
- Shift in focus – IT infrastructure management vs business strategy execution
- Execution – slow time to market

We provide the Cloud Readiness Assessment service, which is an extremely important step in order to plan migrations to the cloud or set up a standby / DR site in the cloud. Whether you move all your workloads to the cloud or adopt a hybrid IT strategy, we will help you understand if cloud adoption is the right strategy for your business. We will conduct feasibility assessment and help you understand the total cost of ownership (TCO). The goal of the Cloud Assessment managed service is to help you build a business and technology strategy so you have a better clarity into:

- How your business will operate
- Possible downtime of migration during migration activity
- Determine Recovery Point Objective (RPO) and Recovery Time Objective (RTO) in case of disaster recovery

When it comes to cloud adoption, it's important to ensure the business and operational processes align and people accept and are willing to embrace change. For this reason, we will help you understand the cloud and how it works, how it can be adopted, and how it can be managed.

As part of this engagement TeraGo will analyze your network, operating systems, and hypervisor. As well as document CPU, disk, and memory utilization at peak and non-peak times. Additionally, we will look at application dependencies, infrastructure resources (i.e. DNS), and performance requirements. Some of this information will be obtained by running a discovery tool on your server, while other information will be gathered manually based on the discussions with you. Our goal is to obtain all of the required information so that we can provide reliable recommendations for the cloud adoption.

Included in this professional services engagement:

- Cloud assessment report
- Cloud assessment recommendations - presentation
- Technical architecture design – current system
- Technical architecture design – future system
- Cloud Total Cost of Ownership (TCO)

Not included:

- Determining the RTO and RPO objectives – this is customer's responsibility

Cloud Migration Services

Doing any kind of migration from one environment to another can be a big challenge and a daunting task. There is a constant struggle to understand which workloads are best suited for cloud migration and which are not. Uncovering the inter-dependencies between applications can also be an overwhelming task, considering a “lift and shift” strategy.

Fundamentally, it is important for you to understand:

- **Impact of migration** – how it applies to your business, what workloads make sense to migrate to the cloud, when should the migration happen, and what downtime can be expect
- **Decision validation** - how to ensure this would be the right decision from business operations, technology, and peoples perspective
- **Business process changes** – what sort of business process changes can be expected when moving to a cloud model
- **Skilled resources** – what kind of resources are needed to help with this effort and the availability of such resources
- **Visibility** – what visibility can be expected into the migration progress across each workload
- **Security** – will the data be transmitted and stored securely in the cloud
- **Testing and validation approach** – how to check if all components are functioning correctly before and after the cutover

With migration services, TeraGo will ensure the identified workloads are properly and smoothly migrated to any of the TeraGo offerings, which may include bare metal, multi-tenant, private and/or public cloud. We take the hybrid IT approach whereby infrastructure, operating systems, VMs, applications, and data is migrated to the right place considering various factors:

- Customer’s business objectives
- Criticality of the workloads
- Network, and application performance requirements
- Business operations and geographic customer presence
- and more

Our approach to migration includes:

- Analyze and create a plan for the migration, which involves creating a runbook that contains workloads, sequence of operations, and the migration schedule
- Performing the migration of one or more VMs at a time
- Cutover to the replica
- Validation to ensure all infrastructure components are functional – data completeness, validation, and performance testing check is customers’ responsibility

Customer Responsibility:

- Functional and performance testing