Design your Azure Infrastructure with Sustainability Best Practices

Author - Nitin Mulchandani

Abstract

- Introduction: We always rush create our Azure landing zone with compliance, security standards and best practices, which is always need of the hour. But at the same time, we need to start thinking on Sustainability
- Abstract: Below are the few areas where we can apply sustainability practices:
  - **Compute**: It can vary from Virtual Machines where we have a larger control where we can implement Right sizing of VM's or implementing features like Auto Shutdown or it can be **PaaS Services** like Azure App service, where we can implement best practices like Autoscaling, Deployment Slots
  - **Storage:** is the heart of many Azure resources like Virtual Machines, Web Apps, Function etc. So, it is always necessary to think of below recommended practices
  - There are other areas like Network, API Management, Serverless compute, Monitoring where we can implement various sustainability best practises to reduce CPU resources

- **Introduction/Overview**

  We always rush create our Azure landing zone with compliance, security standards and best practices, which is always need of the hour. But at the same time, we need to start thinking on Sustainability

  It is always easy to open any cloud portal and spin up a new Virtual machine and same time it is also necessary to understand what happens in the background and how to correct things at design time

  Objective of this whitepaper is not to reduce any security or best practices for any production or non-production environment, but we also have to change our mindset and think in terms of Sustainability, at least starting with non-Production or Developer environment wherever possible.

- **Business Case**

  Typically, now days almost all organization are looking forward or having some quota to give a commitment in implementing sustainability. Whether it is through reducing carbon footprint directly or reducing plastic usage.

  But when it comes to using Cloud providers there are many things which we can do as an Architect to not only right size the Architecture, but also provide best practices

- **Problem Statement**

  Consider an Architect has designed an WAF (Well Architect Framework) which meets the client needs, but there are still improvement areas in terms of sustainability, which can be further tweaked during design time. Because it is easy to deploy things on cloud, but when it comes to changing things, it is always time consuming. So, to avoid the change later, it is always advised to do things correctly during start.

- **Proposed Solution(s)**

  Below are the Areas where we can Apply Sustainability best practices

Compute

- **Virtual Machines**: This is one of the areas where we have a larger control on the infrastructure on the cloud. Below are the few best practices where we can think of **reducing carbon footprint** during design time:

    - **Use B-series VMs**: B-series VM are used for typically idle applications that have sudden usage bursts. B-series VMs use baseline-level CPU power having you paying only for the minimal usage. When there's a sudden burst, CPU power increases and you pay extra for the extra used capacity.

    - **Azure Spot Virtual Machines:** Allows you to take advantage of our unused capacity at a significant cost savings. At any point in time when Azure needs the capacity back, the Azure infrastructure will evict Azure Spot Virtual Machines. Therefore, Azure Spot Virtual Machines are great for workloads that can handle interruptions like batch processing jobs, dev/test environments, large compute workloads, and more

    - **Auto-Shutdown** It is recommended to implement Auto-Shutdown option for your virtual Machine if you are using nonproduction or Dev/Test lab virtual machines, this will enable us to save compute power

    - **DDOS:** Enable Azure DDoS Protection Standard, combined with application design best practices, to provide enhanced DDoS mitigation features to defend against attacks that flood network and compute resources and to avoid unnecessary spike in usage and cost

- **Azure App Service:** This is one of the areas where we are dependent on Azure underlying infrastructure when deploying applications as PaaS. But still there few things which we can manage
    - **Autoscaling:** It is recommended to implement auto scaling for your service plan, this will scale hardware when needed if traffic grows. So, you don't have to procure additional hardware in future
    - **Slots:** It is recommended to enable slots for your App service for non-production environment, which will eliminate the need of having multiple hardware allocation for different environment
    - **CDN:** Azure Content Delivery Network as a global CDN solution to help minimize latency through storing frequently read static data closer to consumers, and to help reduce the network traversal and server load.

    - **Local Cache:** Enable Local Cache in Azure App service to reduce the round trips to backend server

- **Serverless:**
    - **Azure Functions** with consumption-based plan is a great example which are used to replace your classic Web API with Azure Http Trigger functions. Classic web jobs with timer or queue trigger**.** This will remove the need of having

dedicated plan and save lots of CPU usage when not required

- **SQL Database serverless:** The serverless compute tier for single databases in Azure SQL Database is parameterized by a compute autoscaling range and an auto-pause delay. The configuration of these parameters shapes the database performance experience and compute

- **Logic Apps** Is another Azure serverless resource which can be used as background service and can be used to create workflows and orchestrate workflows

Storage

- **Storage Account** is the heart of many Azure resources like Virtual Machines, Web Apps, Function etc. So, it is always necessary to think of below recommended practices
  - **Redundancy:** When spinning up a new Virtual machine Azure recommends a higher storage, which is not required in case of non-production or development environment. So, it is recommended to use LRS (Locally Redundant storage) as Other Redundancy options will need more storage and more CPU, which creates an impact on the environment. LRS replicates your storage account three times within a single data centre in the primary region. LRS provides at least 99.999999999% (11 nines) durability of objects over a given year
  - **Life Cycle Management:** There is an option to enable life cycle management whenever we are using Blobs**.** It is recommended to enable the same for your non-Production environment this which will remove older files and limit hardware space and management on cloud
  - **Azure Storage Access Tiers:** Azure offers different access tiers i.e., Hot, Cool and Archive. Infrequently used data should be stored in **cold** or **offline** archive storage, using less energy.

- **SQL Azure:**
  - **Indexes:** Optimize your SQL server queries by using Indexes this will help to reduce the server usage. SQL Azure also supports Automatic tuning which will constantly monitor your queries and identifies the action that you can perform to improve performance of your workload

  - **Elastic Pools:** Azure SQL Database elastic pools are a simple, cost-effective solution for managing and scaling multiple databases, The databases in an elastic pool are on a **single server and share a set number of resource**s at a set price.
- **Mongo DB:**
  - **Energy efficiency**: Mongo DB provides features like automatic sharding and replica sets, which allow you to distribute data across multiple servers to improve performance and reduce energy consumption
  - **Caching**: Mongo DB supports caching, which can improve performance and reduce energy consumption

- **Compression**: Mongo DB supports data compression, which can reduce the size of the database and improve energy efficiency by reducing the amount of data that needs to be stored and processed

Backup & Logs

- **Geo-Redundant Backup:** Geo-Redundant backup are always need for Ransomware attacks, but should be enabled only for production, which will save CPU which are needed to store replicas. Irrespective of your Cloud resources like Mongo DB, Cosmos DB or Storage
- **Backup duration:** Limit the duration for your logs for e.g., which are stored in Log Analytics if not required

Other Areas

- **Azure API management:**
  - **Implement Rate Limiting:** which can prevent abusive usage of your API's and ensure that they remain available for Legitimate users, this will help in keeping CPU and other resources in control
  - **Caching:** Caching can significantly reduce the number of requests made to your backend servers, which can reduce cost, improve performance and keep CPU and other resources usage in control

- **Network**: Use network security tools with auto-scaling capabilities, such as Azure Firewall Premium or Web Application Firewall on Application Gateway, to enable right-sizing to meet demand without manual intervention and reduce waste of unnecessary resources.

- **Monitor**: For Each resources Azure offers dashboard in the portal to either Self Monitor or implement logs whenever there is spike in errors, CPU usage etc. This will further enable us to further analyse and check the reasons behind CPU spikes if any

- **Application of Solution(s)**

  The current blog best practices are applicable to all Architecture diagrams irrespective of any specific industry

- **Future/Long-Term Focus & Conclusion**

  Cloud continuously evolves so the best practices will keep changing from time to time and it is always good practice keep monitoring things manually as well, this will help us to keep things under control

Appendix – References

https://learn.microsoft.com/en-us/assessments/?mode=pre-assessment&id=a24b1079-29a4-4d22-b678-376e84884f76

https://azure.microsoft.com/en-us/blog/improve-your-energy-and-carbon-efficiency-with-azure-sustainability-guidance/

About Sogeti

Part of the Capgemini Group, Sogeti operates in more than 100 locations globally. Working closely with clients and partners to take full advantage of the opportunities of technology, Sogeti combines agility and speed of implementation to tailor innovative future-focused solutions in Digital Assurance and Testing, Cloud and Cybersecurity, all fueled by AI and automation. With its hands-on 'value in the making' approach and passion for technology, Sogeti helps organizations implement their digital journeys at speed.

Capgemini is a global leader in partnering with companies to transform and manage their business by harnessing the power of technology. The Group is guided everyday by its purpose of unleashing human energy through technology for an inclusive and sustainable future. It is a responsible and diverse organization of over 325,000 team members more than 50 countries. With its strong 55-year heritage and deep industry expertise, Capgemini is trusted by its clients to address the entire breadth of their business needs, from strategy and design to operations, fueled by the fast evolving and innovative world of cloud, data, AI, connectivity, software, digital engineering and platforms. The Group reported in 2021 global revenues of €18 billion. Get The Future You Want | www.capgemini.com

Visit us at www.sogeti.com