HEALTHCARE TRIANGLE

WHITEPAPER

Roadmap to Ensure Successful SAP Migration to Public Cloud

When you choose to migrate your SAP applications into cloud, you will usually come up with reasons like ROI or scalability as to why this migration is required. However, what many often fail to do is, have a proper plan in place as to how and when to migrate their SAP applications to Public cloud. Additionally, questions that need to be answered include:

1. Do you know what data needs to be migrated? Is it confidential, sensitive, and/or critical data?
2. Do you know what OS and DB versions exist on premise? Note: Not all OS and DB versions are supported on cloud.
3. Do you have Infrastructure/SAP architecture diagrams that depict inbound and outbound traffic? These diagrams are critical to study on premise environments.
4. Do you want to migrate the entire data center? Just a few applications? Migration is an opportunity for you to clean up unnecessary data.
5. What types of interfaces are connected with SAP and non-SAP applications? Are you aware of tightly and loosely coupled SAP enterprise business applications?
6. Are you aware of networking and security requirements in cloud? eg: Cloud network and security topologies Regions, VPC, Availability zones, subnets, instances, Network Access Controls, Network Security Groups, Routing tables, Gate ways, VPNs, firewalls etc.
7. Do you have any plans to vacate your on-premise datacenter? If yes, by when?

These are only a few questions out of the many you need to know. In this blog our Healthcare Triangle (HCTI) experts are going to discuss migration techniques that can make your migration go smoothly.

Analyze what you have on-premise

From the beginning, it is important to analyze the SAP applications that are hosted in the data center and understand if you own them. If not, ensure that you inform your hosting vendor well in advance and work with them on a contractual agreement to support your SAP migration to cloud project. It will also be necessary for your team to be certified and trained in cloud and SAP in order to support these migrations. Additionally, start planning your SLAs including backups, disaster restore and recovery, high availability, RTO and RPO. If you don't have certified SAP and Cloud experts on hand, it is always better to start talking to a certified/premium cloud partner like HCTI to get a detailed roadmap for your SAP to Cloud migration plans.

Pilot with what you have

Pilot Selection Criteria: You need to have a variety of applications/OS/DBs for Pilot migration

1. One Linux application without a DB, preferably testing multiple Linux versions (example, RHEL 6, RHEL 7)
2. One Linux application with a DB
3. One SAP application with a DB cluster dependency
4. One Windows application without a DB, preferably testing multiple Windows versions (2008 R2, 2012, and 2016)
5. One Windows application with a DB

Additional criteria wherever applicable:

1. Applications that have no dependency (loosely coupled) with other applications
2. Applications with no shared data storage (SAN/NAS) with other applications
3. Applications with databases less than 1 TB

Technical Migration Planning

On premise - SAP Application versions are not supported on cloud, so you will need to perform patching, upgrade and OS/DB migrations.

There are different methodologies available for migration. Following is an overview of SAP and AWS tools.

Migration Scenario	Source Database	Target Database	Migration Tool or Methodology
Migration of anyDB from other platforms to AWS	anyDB (any non-SAP HANA database such as IBM D82, Oracle Database, or SQL Server)	SAP HANA	1. SAP HANA classical migration 2. SAP DMO with System Move
Migration of SAP HANA from other platforms to AWS	SAP HANA (scale-up and scale-out considerations apply here as well)	SAP HANA	1. SAP HANA backup and restore 2. SAP HANA classical migration (considered a homogeneous system copy in this scenario) 3. SAP HANA HSR 4. SAP HANA HSR with initialization via backup and restore
Migration of SAP HANA from an existing EC2 instance to an EC2 High Memory instance	SAP HANA	SAP HANA	1. Instance resize 2. Amazon Machine Image (AMI) 3. SAP HANA backup and restore 4. SAP HANA HSR

Migration Scenario	Source Database	Target Database	Migration Tool or Methodology
Migration of oracle DB from other platforms toAWS (**OS restrictions**)	Oracle Database	Oracle (Amazon RDS, hardware provisioning, database setup, patching, and backups managed by Amazon) for Non SAP Oracle (Manage yourself) for SAP	1. Cloud Endure 2. Oracle Data Pump Export/Import Utilies 3. AWS Database migration service 4. SWPM (Export/Import) for Enterprise Applications

Oracle-supported operating systems on AWS:
- Oracle Linux
- Red Hat Enterprise Linux
- SUSE Linux Enterprise Server
- Microsoft Windows Server

It might be overwhelming to choose the version you need to use and what tool is required. With HCTI, you can always get the right help at any stage of migration. Whether you plan to move one system at a time or multiple systems in one phase., we would advise you keep it simple and irrespective of approach you follow, plan your downtime based on how much your business can tolerate. However, we would recommend that you move tightly coupled applications together.

Key notes

- Make sure you have solid backups as part of the recovery plan if needed. Store backups to cloud bucket S3 with third party agents BACKINT or store them to FS and run scripts to copy to S3
- Automate tasks for pre-, during- and post-migration
- Have well tested QA before planning for production migration
- Review your HA/DR setup in cloud
- Involve all business units early in order to perform stress test and smoke test to analyze performance on cloud
- Test your business and mission critical applications thoroughly
- Review network, security architecture, and login mechanism in cloud
- Track your billing dashboard

Storage Solutions

Understand storage requirements: Understand the different characteristics, like shareable, file size, cache size, access patterns, latency, throughput, and persistence of data that are required to select the services that fit your workloads, such as S3, EBS, EFS, EC2 instance store.

Evaluate available configuration options: Review various characteristics and configuration options related to storage. Where and how to use IOPS, SSDs, magnetic storage, S3, Glacier, or ephemeral storage to optimize storage space and performance to your work loads.

Best decisions will be based on access patterns: Choose storage systems and configure them by considering how the workloads access data. Make performance improvements, such as choosing caching services or instances that best match access patterns, utilizing optimal key distributions when storing data in S3 or DynamoDB, stripping storage volumes and partitioning data based on system measurements. Increase storage efficiency by choosing object storage such as S3 or block storage such as EBS

Database Solutions

Understand data characteristics: Understand the different characteristics of data in your workloads. Determine if the workload requires transactions and how they interact with data. Use this data to select the best performing database approach for your workloads (e.g., relational database, NoSQL, Datawarehouse, or in-memory storage).

Evaluate available options: Evaluate the services and storage options available as part of the selection process for your workload storage mechanism. Understand how and when to use a given service or system for data storage. Learn about available configuration options that can further optimize database performance or efficiency, such as memory and compute resources, caching, etc.

Collect and record database performance metrics: Use tools, libraries, and systems that record performance metrics related to database performance. For example, measure transactions per second, slow queries, or system latency introduced when accessing the database. Use this data to understand the performance of your database systems.

Choose data storage based on access patterns: Use the access patterns of the workload to decide which services and technologies to use. For example, utilize a relational database for workloads that require transactions, or a key-value store that provides higher throughput, but eventually consistent where applicable.

Optimize data storage based on access patterns and metrics: Use performance characteristics and access patterns that optimize how data is stored or queried to achieve the best possible performance. Measure how optimizations such as indexing, key distribution, data warehouse design, or caching strategies affect system performance or overall efficiency.

Compute Solutions

Evaluate the available compute options: Look at and understand the performance characteristics of the compute-related options available to you. Know how instances, containers, and functions work and what advantages, or disadvantages they bring to your workloads.

Understand the available compute configuration options: Understand how various options complement your workload and which configuration options are best for your system. Examples of these options include instance family, sizes, features (CPU, I/O) function sizes, container instances, single versus multi-tenancy, and so on.

Collect compute-related metrics: One of the best ways to understand how your systems are performing is to record and track the true utilization of various resources. This data can then be leveraged to make more accurate determinations of resource requirements.

Determine the required configuration by high-sizing: Analyze the various performance characteristics of your workload and how these characteristics relate to memory, network, and CPU usage. Use this data when choosing resources that best match your workload's profile. For

example, a memory-intensive workload, such as a database, could be served best by the r-family of instances while a bursting workload may benefit more from an elastic container system such as Amazon Elastic Container Service

Use the available elasticity of resources: The flexibility to expand or reduce your resources dynamically through a variety of mechanism such as AWS Auto Scaling, Elastic Container Service, and AWS Lambda, can help meet changes in demand. Combined with compute-related metrics, a workload can automatically respond to these changes and utilize the optimal set of resources to achieve its goal.

Re-evaluate compute needs based on metrics: Use system-level metrics to identify the behavior and requirements of your workload over time. Evaluate your workload's needs by comparing the available resources with these requirements and make changes to your compute environment to best match your workload's profile. For example, over time a system may be observed to be more memory-intensive than initially thought, so moving to a different instance family or size may improve both performance and efficiency.

To ensure a cost efficient and frictionless migration, it is necessary to have a detailed plan in advance. Reach out to the SAP and Cloud certified HCTI team to learn more about how we can help you be successful with a frictionless and cost-efficient migration of your complete SAP workloads to AWS cloud.

About the Author:

Rajesh Ajmeera, Senior SAP Cloud Architect with 15+ years of SAP, AWS and Azure Infrastructure experienced professional in migrating SAP applications to AWS cloud. He has supported 30+ customers in SAP implementations, projects, and maintenance in multiple roles as an architect, technical lead and manager. He loves technologies and commits to customers to provide best services. He transforms the business of Healthcare Triangle customers, by leveraging emerging technologies across SAP and Non-SAP applications on the Cloud. With a focus on helping large Enterprises get ahead of the Innovation curve, he understands the challenges organizations face and drives forward thinking solutions that help organizations transform their business, by embracing DevOps, Automation and Agile methodology.

HEALTHCARE TRIANGLE

Reinforcing Healthcare Progess™

www.healthcaretriangle.com
(888) 706-0310
(203) 774-3323

For more information, please visit www.healthcaretriangle.com.

©2022 Healthcare Triangle, Inc.
All rights reserved. All other registered trademarks or trademarks are property of their respective owners.

About Healthcare Triangle, Inc.

Healthcare Triangle, Inc. based in Pleasanton, California., reinforces healthcare progress through breakthrough technology and extensive industry knowledge and expertise. We support healthcare providers including hospitals and health systems, payers, and pharma/life sciences organizations in their effort to improve health outcomes. Healthcare Triangle achieves HITRUST Certification for Cloud and Data Platform (CaDP), marketed as CloudEz and DataEz. HITRUST Risk-based, 2-year (r2) Certified status demonstrates to our clients the highest standards for data protection and information security. Healthcare Triangle enables the adoption of new technologies, data enlightenment, business agility, and response to immediate business needs and competitive threats. The highly regulated healthcare and life sciences industries rely on Healthcare Triangle for expertise in digital transformation encompassing the cloud, security and compliance, data lifecycle management, healthcare interoperability, and clinical and business performance optimization. For more information, visit HealthcareTriangle.com.