



Exit Planning for Microsoft Cloud Services

Published: August 2020

© 2018 Microsoft Corporation. All rights reserved. You may copy and use this document for your internal, reference purposes.

Table of Contents

Contents

Table of Contents	2
Introduction & Context.....	3
Exit Planning: Myths & Truths	4
The 7-Step Exit Planning Lifecycle.....	7
Step 1: planning and analysis.....	9
Step 2: risk assessment.....	11
Step 3: decide exit strategy.....	13
Step 4: describe future state.....	15
Step 5: high-level migration plan	16
Step 6: exit plan testing	17
Step 7: plan update	18
Product Specific Guidance	19
Microsoft 365 Exit Planning Considerations.....	19
Azure Exit Planning Considerations	27
Dynamics 365 Customer Engagement Considerations	32
Annex 1: Other Resources	33
Annex 2: Regulation Overview.....	35

Introduction & Context

An increasing number of countries already have or will be publishing/updating guidelines requiring Financial Institutions (FIs) to create exit strategies and supporting plans for their cloud deployments. The goal of the exit plan is to establish a risk mitigation plan for scenarios where, for whatever reason, a FI terminates its relationship with the cloud service provider (CSP).

When it comes to exit planning most of the regulation mandates that a contingency plan may be necessary to exit the outsourcing arrangements with the CSP. **Scoping your exit plans to specific business processes will be foundational since ending your third-party relationship with Microsoft entirely may not be feasible** considering we are both a hyperscale cloud provider and also a major software vendor for on-premise products offering hundreds of mainstream products.

As part of our commitment to transparency, **this document aims to provide you with all the necessary information to support you in creating an exit plan for your Microsoft cloud projects.** It starts off by addressing common misconceptions, next describes a holistic approach for exit planning, goes deeper into each process step and its attention points and finally delivers product-specific guidance for our specific Online platforms. The content is consistent with FSI regulations listed in Annex and includes some viable sample scenarios that you can use for inspiration when building out your own exit plans.

2020 update

We are now publishing a third update to this whitepaper that was originally published in 2018 as a first publication of its kind. In 2020 the overall approach has been adopted by the European Banking Federation (EBF) who use this methodology to publish specific guidance in their paper on testing of exit plans¹.

Some of the changes for 2020 include:

- The exit planning myths and truths were initially added in 2019 and were refreshed in this document
- Added a section on Data Portability
- Updated the Microsoft 365 and Azure-specific guidance
- Added guidance for Dynamics 365
- Added guidance for Microsoft Information Protection
- Links to other resources have been updated across the document
- Regulatory references in Annex 2 were updated

¹ See EBF technical paper: [Cloud Exit Strategy – testing of exit plans](#)

Exit Planning: Myths & Truths

Before discussing the process around exit planning we want to address some Myths and Facts around the topic:



Immediate Business Continuity

The ability to recover from a catastrophic event follows from a resilient solution design without any single-points-of-failure, combined with a tested business continuity and disaster recovery processes for the cloud service.

An exit plan may help to construct a long-term solution but is not an appropriate measure to deal with immediate BCM-type failures².

It is sometimes suggested that the exit plan should provide an immediate recovery from failures by combining use of MS OLS with another third-party (=multi-cloud) or an on-premise (=hybrid) alternative. This would however be a very complex and expensive technical solution that would still not solve non-technical challenges around internal readiness to run and support the environment in the new state. Looking at the risks, it is more logical to ensure the technological solution has high end-to-end resiliency and is absent of any single points of failure. Looking at the cloud components, high resiliency is built into the design of most of our Software-as-a-Service (SaaS) products. Looking at MS Azure Platform and Infrastructure Services (IaaS/PaaS) services, high resiliency can also be achieved but it will be up to the customer to make appropriate design choices for several of these services³. Because of our distributed service design with built-in resiliency, the likelihood that our services would suddenly fail entirely across entire regions or data centers is extremely low. And since failures will always remain a possibility, the FI's BCM-style preparations and exercises must be updated to include the use of cloud⁴. Although there are similarities, the exit plan can never replace these.

Exit Plans document your Exit Strategy

An enterprise needs a strategic alternative for scenarios where the vendor relationship is no longer desirable (continued failures, change in strategy, commercial reasons, regulation changes...)

Each strategic alternative is documented in an exit plan that highlights the current and future state as well as the transitioning process and other factors.

We propose a 7-step approach towards exit planning that works universally for any cloud service.



² For customer guidance on business continuity in the event of Microsoft cloud service disruption, please refer to <https://aka.ms/M365ServiceResilienceGuidance>.

³ The available [Azure SLA](#) also depends upon these choices, see Example 2: Azure Exit Planning Considerations. For instance, for Azure VM the default 99.9% availability SLA can be increased to 99.99% by deploying instances across multiple availability zones.

⁴ Microsoft also prepares for such events as part of our [Enterprise Business Continuity Management \(EBCM\) process](#). We update our customers on the most recent tests on a quarterly basis (link to [our latest quarterly EBCM report](#))

There are various valid reasons that may lead a FI to decide they want to end the cloud outsourcing relationship. These may include a deteriorated vendor relationship or overall dissatisfaction with the received services, a series of repeated SLA breaches, change of internal business strategy or commercial reasons such as failed price negotiations. The exit plan is the ultimate risk mitigating measure providing a long-term solution, but it is not an immediate fix for dealing with a single CSP-related major incident (see Myth 1). Rather, its goal is to help understand the process of ending the vendor relationship, identifying alternatives and assessing the impact on the organization during the transition. This is also our goal: in this paper we provide an end-to-end methodology to ensure an actionable exit plan can be created to deal with these circumstances.



Concentration Risk

Concentration risk is a meta-risk, referring to the extended and potentially systemic impact following a failure.

The best way to manage this is by ensuring cloud solutions have a highly resilient and distributed design preventing any failures to impact the entire service of FI.

An exit plan by itself is not a suitable countermeasure to mitigate concentration risk.

Concentration risk⁵, defined as the systemic risk evolving from a CSP service failure that is used by a large group of FIs across an entire country or region, is not new but has been in existence for many years. Several market data providers & payment service providers offer unique services, with no immediate available alternatives, and while these services have failed at times this has seldomly led to systemic issues. Another misconception has to do with the term “concentration” which is a meta-risk, referring to the increased impact that follows when another risk, such as service outages or provider insolvency, emerges at scale. The most appropriate risk mitigation strategy is to directly address these underlying risks. An exit plan will not solve these challenges.

Risk-Based Exit Planning

Exit planning and testing requirements should be proportionate to the risk involved with the outsourcing.

Most Microsoft cloud services offer a historically proven 99.9% SLA with 24/7 availability; and offer among the highest levels of resiliency in the industry.


Our resiliency design across multiple data centers and regions limits the risk of failures; potentially beyond what is feasible on-premise.



When it comes to exit planning a question that comes up is how the plan can be tested and when testing is appropriate. This specific topic has been explored further in an industry whitepaper on exit plan testing by the

⁵ See also the whitepaper: [Concentration Risk: Perspectives from Microsoft](#)

European Banking Federation (EBF)⁶ where they state that “testing must be balanced against both positive and negative risk factors in order to avoid harming the benefits coming from the cloud outsourcing arrangement”. Building exit plans may carry a high cost disproportionate to the risk. As the EBF paper states: “It may be possible that an institution chooses, specifically, a cloud service specifically to reduce the risk of service failures by deliberately selecting cloud modules that offer a very high resilience or better Service Level Agreement (SLA) compared to on-premises. Similarly, the financial solvency of the CSP, or the product licensing conditions, may result in a lower associated risk when using cloud compared to other IT paradigms. This can reduce the need for exit planning. And finally, exit plan testing may carry inherent risks. Consequently, testing cannot be applied without careful consideration by the financial institution.”.



Vendor Lock-In

Vendor lock-in associated risks are not new, neither are they specific to cloud. Instead it applies to any third-party vendor arrangement also on-premise such as use of licensed hard- and software.

The exit plan allows to end the outsourcing arrangement, but it does not address licensing needs and therefore not resolve vendor lock-in.

Use of cloud will typically reduce vendor lock-in risks due to use of standardized technology that is flexible/portable and therefore easy to transfer.

Vendor lock-in effect emerges when a third-party vendor provides a unique service for which no suitable alternatives are available in the market or on-premise, or when the service does not offer good data portability solutions making it difficult to completely end the vendor relationship. Most FIs have encountered significant vendor lock-in already since many years on-premise, and by switching to a cloud provider the risks associated with the lock-will likely be reduced. The exit plan also will typically not be a solution to deal with lock-in, for instance an Microsoft 365 exit plan is likely to present a scenario to migrate back to the MS on-premise counterparts, ending the cloud outsourcing arrangement but not the vendor relationship in its entirety.

Planning for a new future

When things go wrong the exit plan delivers an enterprise with real-life, feasible alternatives beyond the use of cloud.

Exit planning may also require you to modify or end business processes that depend upon current use of cloud.

Planning alternatives and reviewing or approving these with internal stakeholders is best done when not yet under stress.



⁶ See EBF technical paper: [Cloud Exit Strategy – testing of exit plans](#)

Not every cloud service can be easily transferred to an alternative solution. Sometimes the cloud service is quite unique and the technological alternatives may be hard to define, requiring custom software development and/or complex migrations of data and business logic. In these cases, sometimes the preferable solution involves changing the business process itself so that these costs and complexities can be avoided. This process of identifying alternatives must be done together with many stakeholders, including business, and is best done long in advance of any actual problems that might trigger an exit.

Is exit planning necessary?

Sometimes. From a regulatory perspective this will depend upon the region in the world where the cloud service is being deployed and the applicable regulation within that region. For instance, in Europe, exit planning is only necessary for critical or important functions being deployed in the cloud. The principle of proportionality is also applied, not every deployment may require an exit plan.

But even if you are deploying only in regions that have not issued specific guidance, it may still be wise to develop an exit plan to limit associated risks to the enterprise. Also, it is not unlikely that more countries will introduce exit planning requirements into their outsourcing regulations.

Exit plans are also best reviewed at regular times. When our FI customers deploy cloud services we see the usage of our services gradually growing over the years. The amount of services being used and the amount of data stored can vary greatly over time and since these elements are foundational to the exit planning process, setting up a scheduled exit planning review and update process is essential for success.

How much effort should go into the exit plan? How should we test it?

Creating an actionable and tested exit plan – depending on how it is defined – can become an extremely expensive undertaking, possibly to such an extreme that it subdues your original business case for use of cloud. However, most regulators will allow you to take a more reasonable, risk-based approach towards building and testing exit plans.

Exit plans must be balanced against both positive and negative risk factors in order to avoid harming the benefits coming from the cloud outsourcing arrangement in the first place. Testing may carry high costs, disproportionate to the risk of the service failing in the first place. It can also be possible that an institution chooses a cloud service specifically to reduce the risk of service failures by deliberately selecting cloud modules that offer a very high resiliency or better Service Level Agreement (SLA) compared to on-premise. Similarly, the financial solvency of the CSP, or the product licensing conditions may result in a lower associated risk when using cloud than using other IT paradigms. These factors may diminish or reduce the need for exit planning. And finally, exit plan testing may in themselves carry inherent risks. Consequently, testing cannot be applied without careful consideration by the financial institution.

The 7-Step Exit Planning Lifecycle

To help our FI customers meet regulatory requirements around exit planning we highlight in this document a template approach towards exit planning, we included some of the applicable regulations and finally we highlight a few high-level migration scenarios specifically for exiting our MS Online Services.

Below is an overview of our 7-step exit planning lifecycle:

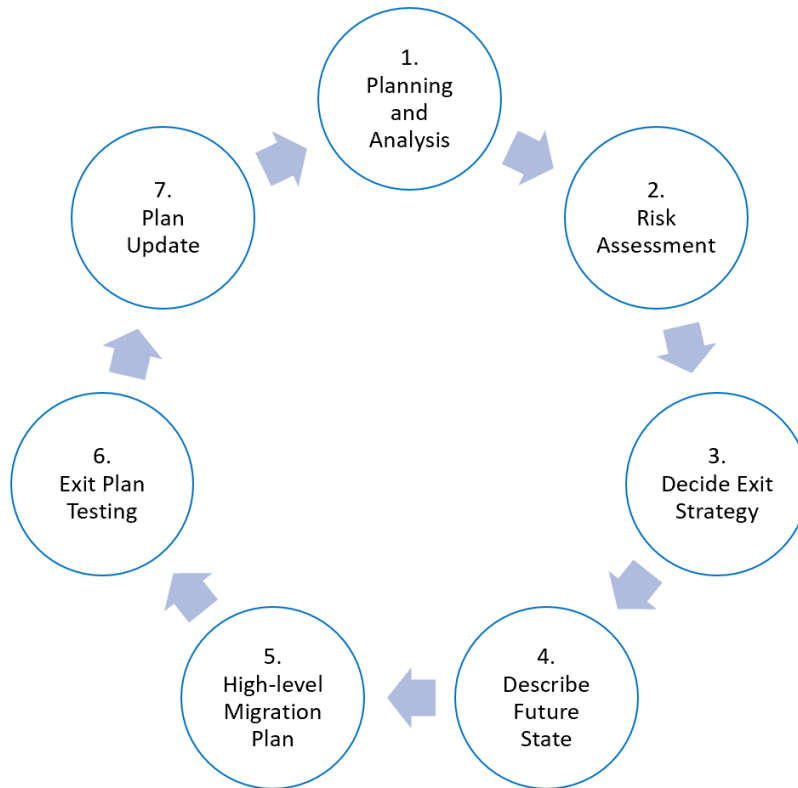


Figure 1 Creating and maintaining an exit plan in 7 steps

These 7 steps reflect a holistic approach towards establishing an actionable exit strategy. This approach is aligned with the European Banking Federations⁷ guidance on exit planning as published in 2020⁸.

Before walking through this lifecycle, we want to clarify the difference between exit strategies and exit plan:

- **Exit Strategy:** A high-level description of an institutions' ultimate risk mitigation strategy when dealing with a failing cloud provider or when terminating the outsourcing. This might include exit and transition of outsourced functions and data to an alternative provider (in part or completely), the return of these functions on-premises, or even discontinuation of the process.
- **Exit Plan:** A high-level description on how to implement this exit strategy, including a description of all its phases, involved roles and responsibilities and various plan features. A plan is to be enacted in case of predefined events following a long-term strategy approach. It does not include short-term incident management, since the business implications of enacting an exit plan can be considered severe. The exit plan ensures business continuity in case of the predefined events, aiming at response times appropriate to the severity of the triggering event.

Our iterative approach fits very well within a phased deployment model where more and more of the cloud services are gradually being consumed over time, which is the most common usage pattern we see across industry.

⁷ The European Banking Federation represents the European banking sector by uniting 32 national banking associations in Europe that together represent some 3,500 banks – large and small, wholesale and retail, local and international – employing about two million people. Regular forum meetings are organized to address the use of cloud within the banking sector.

⁸ See EBF technical paper: [Cloud Exit Strategy – testing of exit plans](#)

Initially exiting the cloud is often reasonably straightforward, and the initial exit plan can also be concise and simple in nature. However, as time continues hybrid deployments may transition into cloud-only environments, more data is being stored online, less in-house skills become available and the exit plans must be further matured to ensure they still reflects the new reality.

Step 1: planning and analysis

Goal of this step is two-fold: make sure all preparations are made including identifying stakeholders and ensure the current use of resources both in cloud as well as on-premise is understood (=the as-is configuration of the service in scope of the exit planning process).

Setting up an Exit Plan Coordination Team

Writing a well-founded exit plan requires the involvement of different stakeholders across the organization. Key tasks include identifying who are the different actors engaged in your internal organization, agreeing on who is responsible for what and deciding on common criteria for the plan.

For this we recommend setting up an **Exit Plan Coordination Team** which is a group of individuals that represent different stakeholders in the entire sourcing & exit process. This team will meet during the creation of the initial exit plan, and thereafter reconvene at least yearly to update this plan and amongst other items, discuss who is responsible for each task. A clear owner should also be appointed to manage this process.

We recommend linking key individuals to key exit plan tasks and responsibilities in a responsibility assignment matrix or RACI table⁹ which can be reused in several of the later steps of this plan. By defining the RACI participants upfront, you can ensure everyone involved in the development of the plan is informed from the start and kept updated throughout the process.

Below are some possible participants in the Exit Plan Coordination Team:

Role or function	Name and Contact Information	Role Description
IT Leadership (CIO and/or CTO)	...	Accountable and responsible for delivering the business service and chairs the Exit Plan Coordination Team. The IT leadership must ensure the availability of budget and resources to maintain the exit plan for the service. We also recommend IT leadership to setup vendor and internal SLA monitoring, as discussed in step 2, because they are responsible to deliver the internal SLA to the business and will likely play a key role in the events leading up to the exit plan being triggered. An exit decision due to service failures should be based upon sufficiently reliable measurement data.
Business owner(s)	...	Key business representative(s). They represent the business(es) that would be impacted when an exit is triggered. It is important that these function(s) have delegate decision power by other business lines that may not be directly represented, as is often the case with a broad service such as M365. The business determines the required SLA for the service together with IT leadership. Another key role for business is to evaluate if

⁹ Responsibility Assignment Matrix or RACI matrix (more [info on Wikipedia](#)).

		the future solution still meets all business requirements and if any identified gaps in functionality are acceptable.
Cloud Architect / Technology Architect	...	Their focus will be on documenting the current state in step 1, for proposing one or more solutions in step 3 and finally also for documenting the future state in step 4. They are also consulted in the migration planning in step 5.
Procurement and Vendor Management	...	Owns the contractual relationship with Microsoft. They are responsible for the contractual exit out of the service and for the procurement of new solutions where applicable. In addition, they support IT leadership during escalation processes in the period leading up to the triggering of the exit plan. Moreover, they may also be involved in upstaffing and contracting efforts during the exit plan execution.
Partners, sourcing providers	...	(optional) In some cases a third-party will get involved in exit planning, and agreements on the amount of work that must be delivered can be agreed upfront in a statement-of-work (SOW) linked to the exit plan. These partners must agree upon the workload estimates & feasibility of the plan.
Legal and Compliance	...	(optional) Supports IT leadership during the escalation process and as a stakeholder is involved during the assessment performing a high-level screening of the future solution to determine if all compliance requirements can still be met.
IT Security	...	Primary involvement is to perform a high-level assessment of any future solution, validating if security requirements will still be met. Moreover, the security of the exit plan itself must also be validated.
HR & Communication Departments (=business)	...	(optional) These are some of the business functions that are more directly involved in the changes following an exit, and we recommend adding them into the exit plan coordination team for this reason. HR may play an important role in the required IT staffing or role changes, while the communications department will manage the communication plan.
Project Office	...	(optional) Exiting a cloud service is a not to be underestimated project, likely both urgent and of high importance. It is essential that the project office delivers the right profile(s) to manage the exit successfully. They are often responsible to identify and deliver the right profiles to manage the exit plan execution.
Finance	...	(optional) Involved in assessing the financial implications when the exit plan is triggered, which will usually lead to unforeseen and unbudgeted expenses for the organization. They will also assess any future providers from a financial perspective.

The names and contact information of the *Exit Plan Coordination Team* should be documented as part of this plan (=requirement in specific country regulations, for instance in Japan).

Determine plan governance cycle

Another item that falls outside the plan itself is determining how it will be governed and maintained. Questions that must be answered include:

- Who will approve the exit plan? Exiting a cloud service often has implications across many business lines and divisions, and not each division may agree that the service is degraded to an unacceptable level

mandating an exit (and the associated overhead). We recommend determining upfront which committee or decision body is authoritative to make such a decision.

- How will the plan be maintained? What is the update frequency of the plan? Who is responsible for creating updates? What form of testing – if any – is appropriate for validating the plan contents (e.g. execution of an annual table-top test with all stakeholders).

When it comes to the evaluation process of triggering a possible exit scenario it will be necessary to have these preparation steps in place in order to avoid confusion over authority, internal disputes or chaos.

Defining Key Success Criteria

The objectives of the exit plan, which reflect the final conditions that must be met for this plan to be executed successfully, should be determined upfront. We recommend determining some Key Success Criteria following the guidance from the European Banking Authority (EBA). These criteria can be defined in terms of cost, availability, resilience, functionality or other terms (ideally measurable criteria are used).

A clear statement on the level of expected depth for the plan may also be helpful. Is it acceptable to setup a plan that highlights principle steps only, or does the organization desires to support this upfront with specific commitments on financial, people or technology resources? What are the quality criteria that must apply?

Listing of services currently running in the cloud

When planning for a service exit, a first step is to get a good understanding of the current state of your cloud deployments. There are several questions to be answered:

- Which services are we using today? Are these in production (vs. being piloted)?
- How many users are using the service?
- How much data sits in the cloud today?
- How is the usage different across different divisions in my organizations?
- Other specific requirements...

A high-level overview is provided of which cloud services are effectively in scope for the exit plan. An overview of all Microsoft Online Services can be found in our [licensing terms](#). Only a few of these will be effectively in use within your organization.

When describing the current state, the list of cloud services should be augmented with short business service description and other useful info such as the amount of data stored online. In this section it is also useful to include information like executed business impact assessment (BIA).

Examples for Microsoft 365 (M365) and Azure can be found in the Examples sections.

Step 2: risk assessment

In this phase we describe at a strategic level which threat scenario could ultimately lead to an exit being triggered and highlight the high-level exit strategy that would need to be applied for each service so that the impact becomes clear. Different approaches may be taken here.

Identifying Key Risk Indicators

Part of the threat analysis involves brainstorming over the different threat or risk events to the local business that would ultimately drive the decision to exit the service (key risk indicators), and document these. Taking into consideration exit planning cannot provide an immediate solution for imminent failures, we focus on structural issues and risks that may trigger an exit.

Impact Analysis vs. Exit Plan

The potential business impact is different for different possible threat scenarios (see key risk indicators), and it should be determined which of the threats receives focus in the exit plan. You may also want to set high-level objectives upfront as to the desirable timeframe during which an exit should be achievable.

The most drastic scenario will involve a full exit, but in organizations where mergers and acquisitions are common it may be desirable to also consider carve-out scenarios that support the exit within a single subsidiary or a single country as an annex to the baseline (full exit) plan.

Service-Related Exit Triggers

An exit of a cloud service will cost the organization substantial time and effort and is often a measure of a last resort (an ultimate risk mitigation). It is not always obvious where the precise point lies where one considers the service to be so heavily degraded that it mandates these investments.

Service-related exit triggers could include:

- SLA breaches spanning multiple months, resulting in material financial losses
- Major breach of trust due to a major service incident, for instance a multi-day outage

We advise brainstorming upfront on the maximum levels of unavailability after which the plan would be triggered. The intent of such a trigger is to act only as a guideline, where the formal decision to exit the service will need to be taken on a case-by-case basis by the Exit Plan Coordination Team or similar body.

Important note: We recommend to setup an internal process to ensure effective SLA measurements are also available within your organization. Note that Microsoft does not measure or publish detailed SLA performance numbers, only [aggregate numbers for the entire service](#) are made available at a high level. Therefore, we recommend to setup measurements of SLA performance against your own tenant, allowing you to get a service reporting which is fully representative for the end users in your organization. If vendor SLA monitoring is also required, another option can be to rely upon a third-party service for this.

Contractual Exit Triggers

For Microsoft Cloud Services there are a number of contractual provisions which may trigger an exit documented in the [Online Services Terms](#) (OST) which may result in a (forced) service exit. An exit can also be triggered by Microsoft, for instance due to changes in regulation within specific countries preventing us from continuing to deliver the service in that location; or an exit may be triggered by your organization due to concerns over specific service changes (for instance due to an unacceptable change in Microsoft's list of subprocessors). And contracts may be terminated both by customer and CSP in relation to a major security incident

Other Exit Triggers

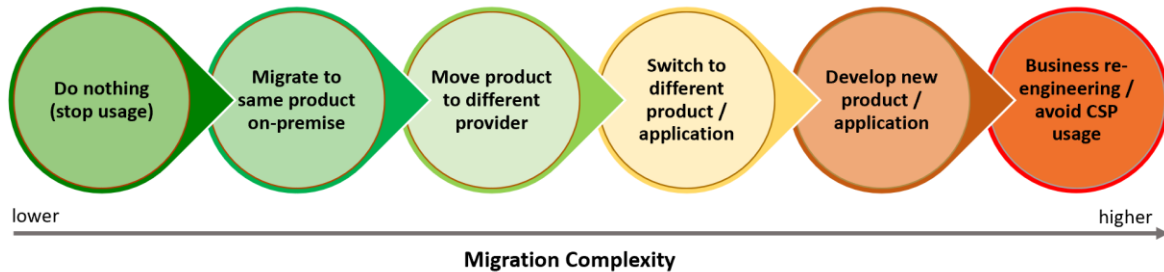
Finally, there also exist a myriad of other events of change that may trigger an exit which may be hard to predict and therefore cannot easily be anticipated in the exit planning process. For instance:

- Changes in vision/strategy
- Commercial reasons such as a price increase or a competing offer by a different provider
- Business divestments, for instance when a subsidiary or activity is sold off
- Geo-political crisis
- Disputes or ongoing litigation
- Other

Step 3: decide exit strategy

Evaluating Future Solutions

The general idea is to list up the different possible exit scenarios at a high level, and weighing their pros and cons.



Evaluations must drill down to the different cloud subservices in use since not every alternative may be viable for each application.

Different exit scenarios are available for different applications, and we recommend listing each of these in this section together with their strengths and weaknesses. A preferred exit scenario is selected for each application, and only the preferred scenario is taken forward in the next steps of this exit strategy document.

Next to the technical feasibility, complexity and cost that distinguish the different available options, attention must also be given to minimize the user and business impact of a cloud exit. Assessing this impact is very important because this may become a determining factor in next steps of this exit plan.

Dealing with cloud-native applications

Exit planning becomes particularly important for any cloud (sub-)services that don't have an on-premise counterpart. In these cases, the FSI must revisit its internal use cases and consider on-premise or third-party alternatives that offer similar functionality. Or – alternatively – discontinuing the service altogether during an exit may also be considered for non-critical services.

Deciding a Sourcing Strategy

An exit plan can become a very complex and work intensive undertaking in certain scenarios, and your local organization may decide to source externally for the execution of the exit plan. If this is the case, potential sourcing partners may be listed as part of the exit plan. If the organization has a preferred sourcing partner, it may be

beneficial to involve the partner also in the exit planning process asking them to validate the future state & migration planning estimates.

In some cases, especially for material deployments which have reached a point where reversibility has reached a significant level of complexity, it may also be beneficial to already prepare a statement of work (SOW) tied to the exit process. This will increase the level of confidence that the exit plan can be executed as planned.

Prioritization: Speed vs. Quality

As with any large project it is important to be clear on priorities when planning an exit. For instance, if an exit must happen under stress or duress for some reason, an organization may want to prioritize its efforts in a way that an alternative business process can be setup robustly and quickly, potentially sacrificing some less important requirements. For instance, the solution may choose on ensuring an alternative solution is established that works but disregards older archive data that cannot be migrated out of the solution quickly, or disregards certain subprocesses that are not essential.

Determining Migration Strategies

One of the major challenges involved in executing the exit plan is to determine how business logic, data and metadata can be migrated to the future solution:

- Even in seemingly straightforward scenarios the interoperability between the current cloud system and the future alternative may have limitations, especially if different product vendors are involved.
- The same is true for data & metadata portability, not every piece of data within the current environment may be transferable, and for data that is transferred data interoperability issues may arise
- Finally, financial institutions – for good reasons – set high security and compliance requirements for cloud solutions, but these requirement may become impose additional challenges when actioning on the exit plan.

For instance, containerized environments may seem good candidates for an easy transfer on the surface, but the data stored within them may be subject to legal data residency requirements within a specific geography, and/or may have requirements to stay protected by FIPS-140 compliant hardware encryption keys stored in Hardware Security Modules (HSMs), both of which can further complicate the exit planning process.

Another major challenge when moving back on-premise is the migration of large sets of data across a network with only a limited amount of unused bandwidth. For software services we may also throttle the amount of resources and bandwidth that can be consumed by a single user, which may also restrict the speed at which data can be transferred across the Internet.

Depending upon the size of the data set being consumed in the cloud (as identified in Step 1), different approaches for exiting the data out of the cloud may be considered.

Rather than moving all the data across the network, in specific cases a distinction could be made between *live or hot data* and *cold data* and different data migration strategies could be applied to each category to limit the amount of data being copied across the Internet:

- **Live (hot) data** meets 2 criteria: (i) this data must remain accessible to end users with minimal downtime, and (ii) copies are not available upfront on-premise to support exit purposes. If we take the example of Microsoft 365 Exchange online, this would be the non-archived data in the live mailbox of a user. This is the data for which finding the right migration path becomes critical!

- **Cold data** is any other data: data in the cloud that is not commonly used/changes & data that is also available on-premise. The idea is that this data either already exists somewhere else or can be shipped to a new solution without time pressure, rather than having to be copied over the network. In our M365 example below, this could be all data in archive mailboxes that a user does not need to have immediate read-write access to.

Data Portability

Data portability ensures you can transfer your data to another solution, which is a foundational necessity in order to establish a working exit plan. Several aspects have to be considered such as the data being accessible for transfer; use of in a machine-readable data format that can be understood by receiving applications correctly for instance through use of documented application programming interfaces (APIs); applicable metadata that is required for the data to be meaningful; data security measures such as use of data encryption and related key management/protection and finally data compliance requirements where laws or regulations prohibit transfer to a certain location/region.

There are some good standards documents that dive deeper into this topic:

- [ISO/IEC 19941 on Cloud Interoperability and Portability](#) highlights in more detail these different facets that must be considered when migrating data between different environments.
- Draft [ISO/IEC 19944 on Data Flows, Categories and Use](#) highlights a possible data taxonomy that can help in planning data migration activities.

Data portability is also embedded within EU regulation¹⁰ and led to the establishment of the [SWIPO](#) initiative. SWIPO (Switching Cloud Providers and Porting Data), is a multi-stakeholder group facilitated by the European Commission, in order to develop voluntary Codes of Conduct for the proper application of the EU Free Flow of Non-Personal Data Regulation / Article 6 "Porting of Data".

You will find on the SWIPO website links to two different codes of conduct for Infrastructure as a Service (IaaS) and Software as a Service (SaaS), and Microsoft is part of this initiative together with other large cloud providers.

The actions taken under SWIPO will however be beneficial to all Microsoft customers worldwide in support of the creation of exit plans, not just in Europe.

Step 4: describe future state

Exiting a cloud service in its essence is a migration of data and functionalities across two platforms. To plan the migration activities that must be undertaken (=described in step 5), a good understanding both the current (as-is, described in step 1) and future (to-be) states must exist.

In this section we drill down deeper on the future state of service (post plan execution), deepening the principal decisions that were made in step 3.

There are several areas to cover such as:

1. General Considerations: what do we assume when documenting the future state
2. Technology: which hardware to purchase, licenses to use, software setup etc.
3. People & Support Requirements: what will the future support organization look like, and

¹⁰ See [EU Free Flow of Non-Personal Data Regulation](#), Article 6 "Porting of Data".

4. Impact on End Users and Business: specific attention points we recommend giving specific attention

We'll explain most of these by coming back to the examples used throughout the document, since these scenarios differ for each deployment.

People & Support Requirements

High-level description of the supporting environment for the future state. Which roles, skills and support resources are necessary? How will the future solution integrate into the IT management plane?

Impact on End Users and Business

The objective of the exit plan must be to minimize visible impact following from a rollback process. To establish this, services must be established on-premise and end-user workstations must be redirected to the new environment ideally without too much impact in terms of client software changes.

However, as indicated before, some scenarios may require users to change their habits and/or learn to use new technology solutions. If this is significant the management of change must also be planned and documented in the exit plan.

Step 5: high-level migration plan

The high-level migration plan itself can now be created. This plan walks through the different steps and milestones required to effectively establish an exit from the cloud service, gradually reaching the future state as described above. A convenient form would be to draft some key elements of a (future) project memorandum, in many ways similar to your onboarding project memorandum¹¹ but in the opposite direction.

A note on the expected level of depth: we suggest a common-sense, risk-based and balanced approach in line with the EBF guidance paper. The overall exit plan, and this phase of the high-level migration planning, can be of any size from a short summary of steps to a detailed project plan with some higher-level technical design elements. In either case, **we are unaware of any regulatory requirements that require FI's to provide detailed technical and operational procedures describing how exactly one exits the service (at a very deep technical level)**. We believe this will also be unrealistic in modern-day environments where DevOps and Agile techniques bring constant changes to the production environment. Such an approach would likely require significant investments, disproportionately large compared to the risk that such a plan would ever need to be invoked, and the benefit would also be short lived unless the document is maintained on a continuous basis (adding to the overall cost of having and maintaining this exit plan).

Creating a Migration Plan



Preparation



Planning



Build solution



Network
upgrades



Pilot
migration



Velocity
migration



Cold data
migration



Wrap-up

¹¹ The onboarding project memorandum & technical design documents may be excellent resources to establish initial exit strategy documentation.

Above are some sample milestones you may want to include in your migration plan with some level of detail. For each of the milestones we suggest a brainstorm session with all stakeholders on:

- The different steps that exist within each phase
- The amount of effort required (e.g. in man-days)
- The amount of throughput time needed for the entire phase (or for each sequential step)

This will give you a realistic feeling about the budget and throughput time needed for an exit. A table-top planning exercise and [planning poker](#) techniques with the *Exit Plan Coordination Team* can help in establishing realistic estimates for the exit plan.

Technical Migration Highlights

For core parts of the migration you may want to include some technical highlights, but we recommend keeping this limited due to the high effort this would require.

In Annex and in our Examples, you will find several links that can be helpful in defining high- and low-level exit scenarios for your specific context. It is also possible to involve [Microsoft Enterprise Services](#) to support your organization in creating a technical Microsoft 365 Rollback Plan tailored to your organizations need. Note that Microsoft Enterprise Services can also offer enterprise support services for the plan execution.

IT Training & Education

There may also be internal limitations to the amount of expertise available to manage the underlying infrastructure and application environment on-premise once the rollback plan has been invoked. In case of an exit, additional competence must often be acquired externally. Contact your account team or [Microsoft Enterprise Services](#) for a statement of work to deliver additional support as needed.

In addition, your first-line help desk workers must also be trained to handle a possible increase in calls as a result of the migration back on-premise. Enough time and resources should be foreseen in the exit plan to ensure this.

Lastly, the rollback may affect how the service is delivered to end users, requiring some education of end users (management of change).

Update list of Owners (Key Contact Persons) with Contact Info

Some national regulations require organizations to maintain a list of key contact persons with their up-to-date contact information so that these can be reached as part of the contingency plan (e.g.: Japan). Regardless whether your organization has presence in these countries, we recommend as a general guideline to maintain contact information as part of this plan. We also recommend focusing on identifying the owners for the different migration plan steps, considering them as key contact persons for each step of the plan.

A table is generated as follows:

Role or function	Responsible Individual	Contact Information (telephone & email)
...
...

Step 6: exit plan testing

After having gone over all the previous steps a draft plan now exists. As a final step we recommend performing an exit plan validation in an exercise like a table-top business continuity testing process. During such a process, the key stakeholders are brought together in a large meeting room and the team walks over the plan to discuss the accuracy of all estimates and assumptions that have been made.

Some questions that can be discussed during such a meeting include:

- Have we identified all the necessary stakeholders? Do we need to include others in the coordination team? Are the accountabilities, responsibilities sufficiently clear? Update of owners.
- Validate (update) and walk through the assumptions of the current state. How much data is consumed by the cloud service? Are any new cloud features or functions in use, and what does this mean for the exit plan?
- Threat analysis review. Are there any new threats that can be identified? Do we need to work out a new scenario? Evaluation of exit triggers. Are these still appropriate?
- Is the future scenario still the most appropriate? Are the assumptions still valid? Does it cover a plan for any newly added functionalities?
- Walk-through of the future state and migration plan & updating of any outdated figures
- Sign off on the plan update

Step 7: plan update

As a final step, each year, the plan must be revised to check if all assumptions are still accurate. This is also a good moment to revisit the amount of cloud services being used and assess the volume of data being stored in the cloud.

Product Specific Guidance

So far, the guidance we provided has been universally applicable, but in the following examples we want to guide you through some practical tips and/or examples that are specific to the deployment of Microsoft 365, Azure and Dynamics 365.

The most exhaustive example below works through all 7 steps of the exit planning process in context of M365. For Azure and Dynamics 365 only the product specific items have been highlighted.

Microsoft 365 Exit Planning Considerations

Microsoft 365 is a Software-as-a-Service (SAAS) platform that offers a set of different services, and enterprises can choose to deploy these in a varying number of ways. Many financial institutions will also start their deployments in a limited way, for instance only deploying a core set of products like Exchange, SharePoint and Teams but over time gradually consume also other services. Another possibility is that an enterprise starts with a hybrid setup keeping some of the data and infrastructure on-premise, and gradually move more and more workloads into the cloud over time.

Rather than to talk to all these scenarios individually, we documented a hypothetical scenario for a small financial institution that has migrated some of its workload into the cloud (non-hybrid). In step 1 we describe this entirely hypothetical yet also very realistic scenario, walking through the different exit planning phases in following steps.

Step 1: planning and analysis

Figure 2 represents the current deployment state of a hypothetical company that has deployed, or intends to deploy, M365 and needs an exit plan. The current state must also indicate the date at which the info has been gathered so that during later planning phases correct assumptions are made on the amount of data in use.

Online M365 Tenant status as per 1 January 2018			
M365 Service	Business Function	In use?	Volume/comments
Exchange Online	Personal mailbox, calendar & contacts	Y	1200 Mailboxes; total 3 TB
	Shared mailboxes	Y	300 Mailboxes, 1.5 TB
	Email archives	Y	1500 Mailboxes, 45 TB
	Mail flow security EOP, ATP	Y	
	MX in/out mail services	Y	
Azure AD (AAD)	AAD User Provisioning	Y	1000 users, synced via AAD Connect
	Federated authentication	Y	ADFS
InTune (EMS)	Mobile Device Management	N	
Azure Information Protection (AIP)	Document-level encryption of documents classified as being secret	Y	Estimated a few hundred emails and documents
SharePoint Online	Intranet publishing (websites)	Y	45 site collections; 500 GB of data
	Document collaboration	Y	4 TB
OneDrive	Local file storage & sync	Y	6 TB of data
	Personal document collaboration	Y	
Skype Online	Online meetings & calls	N	Using S4B on-premise deployment
Teams	Online collaboration	Y	Pilot within IT, 150 users

Online M365 Tenant status as per 1 January 2018			
M365 Service	Business Function	In use?	Volume/comments
			Total: 60 Tb of data

Figure 2 Sample table of services in use + indication of data storage volumes

Assessing Business Impact

List up the internally available information such as the business impact assessment (BIA) which assesses the ultimate business impact of a service failure for each business process across the enterprise. Next, map the cloud services and relevant on-premise infrastructure onto this BIA, creating an application and infrastructure level view on the business criticality for each of the cloud services. For M365 we recommend that the criticality is described at the level of the individual M365 subservices, as each service may not have the same level of business criticality.

The outcome would be a list of M365 cloud subservices with their business criticality expressed both in terms of SLA requirements and maximum recovery time expectations, in accordance to this BIA assessment.

For instance, using the example above, you get the following (cloud components only):

M365 Service	Business functions (from BIA)	MS SLA	E2E SLA for the business
AAD Provisioning	Provisioning subservice that spans multiple business functions: <ul style="list-style-type: none"> - HR joiners-movers-leavers processes - Adding & removing other users (contract staff) 	99.9%	95%
AAD Authentication	Authentication subservice that spans multiple business functions: user authentication (login)	99.9%	99%
Exchange Online	Email services	99.9%	98%
SharePoint Online	Intranet & internal websites	99.9%	99%
OneDrive	Personal file storage	99.9%	98%
Teams	Online internal collaboration (pilot within IT)	99.9%	N/A (PRD 98%)

Figure 3 RTO requirements per application resulting from BIA

Your internal end-to-end SLAs as it is offered to business users must always be lower than or equal to what Microsoft is offering for the service, since this includes any potential downtime to other components such as the network connectivity and on-premise infrastructure (e.g. ADFS, ADFS proxies, AD, AAD Sync..).

Step 2: risk assessment

As discussed in the introduction of this paper (myth 1), exit planning involves the ultimate risk mitigation but should never be your primary measure in dealing with CSP service failures.

Since one of the reasons for invoking an exit plan may be due to continued failures in the service, you may want to consider defining some exit plan triggers upfront that put clear lower limits on when an exit may become appropriate (due to a loss of trust in the provider). For instance, you could argue that repeated SLA breaches below 95% are unacceptable and would lead to a breach of trust in the provider, triggering an exit out of the service. Another trigger could be a multi-day failure alongside some smaller issues, again making it very difficult to maintain trust in the provider.

Step 3: decide exit strategy

Below are some considerations for organizations that are migrating back on-premise:

- **Exchange, SharePoint Online & OneDrive** on-premise alternatives offer an easy 1-to-1 migration by executing a reverse onboarding process. Still, caution must be taken because a small set of cloud functions don't have an on-premise counterpart anymore and may require third-party products to replace them on-premise (e.g.: some security and compliance features are cloud-only solutions and don't have an on-premise counterpart).
- **Identity migration** depends upon the deployment model, but generally Azure AD and on-premise AD are very compatible and different migration scenarios are offered. The caveat will be – again – that all used features must be evaluated to see if and how these can be deployed on-premise.
- **Azure Information Protection:** move the information back on premise, relying upon AD-RMS encryption as an alternative to the Microsoft Managed Keys (MMK) used in the cloud instance. A description of the migration path for AIP is documented in Annex 1.
- **Teams, Yammer & several other apps** that are part of the M365 suite do not have any alternative on-premise and require users to consider alternative products or approaches to satisfy their use cases. Without being specific on how to migrate these applications, the drawing above highlights some of the alternative approaches that one may consider (see also below – dealing with cloud-native applications).

Microsoft Teams & Yammer exit planning considerations

As of time of writing both Yammer and Teams are storing content data within your tenant environment, using the underlying MS Exchange and SharePoint subservices. This data can be migrated easily to an on-premise SharePoint environment. The main challenge lies with the social functions and metadata around these services that will require alternative solutions. Possible alternatives that can capture some of these functionalities on-premise include Exchange lists and public folders or to rely upon SharePoint on-premise functionalities combined with a third-party solution to cover social collaboration requirements gaps.

Office Apps planning considerations

Your locally installed Office apps will either involve a traditional on-premise Office deployment (e.g. Office 2019), or alternatively your organization will have deployed [Office ProPlus](#) which is licensed to the cloud service. Next to switching the license and deployment models as needed, some [intelligent services](#) may be in use that are exclusively available as a cloud service (for instance: translation of text & accessibility features). These requirements must also be considered as part of the exit planning.

Live (hot) data vs. Cold data migration

While live data migrations are almost always done over the network, there exist several alternatives in dealing with cold data:

Keep cold data in O365	Migration across the network	Own process setup & data disk shipping
Scenario focuses on Exchange Exchange (cold) data moved to archive Information access using web clients OneDrive data already locally available & synced to new solution Smaller Exchange hybrid migration for hot data across the network	All data is exported over the network Priority migration of hot data Cold data: use long transition period	Develop data migration factory on Azure IaaS Supported by MCS or MS Partner Setup VMs with O365 & 3rd SharePoint tools Connect Azure Blob Storage to VM on Azure Export Cold Data on Azure Blob Storage Ship hard drives to MS DC & transfer data Return data hard drives to on-premise DC Move data back into on-premise applications

Figure 4 Data migration scenarios

Bringing it all together

Following the example throughout the document, we calculated the amounts of live vs. cold data that would need to be migrated. The archiving policy was setup to archive all Exchange email information older than 6 months, maximizing the amount of cold data. This provides us with the following sample view:

Business Function	Live Data	Cold Data	Remarks
Personal mailbox, calendar & contacts	3 TB		1200 mailboxes
Shared mailboxes	1.5 TB		300 mailboxes
Email archives		45 TB	Everything older than 6 months
SP publishing (websites)		0.5 TB	Source data was already available on-premise
SP collaboration (collaboration)	4 TB		Live data because it is not available on-premise
Personal OneDrive files	6 TB		We will treat this as hot data due to the limited size (but this could be treated as cold data also)
Online collaboration	0.1 TB		Teams data & metadata
Totals	14.6 TB	45.5 TB	

The available network bandwidth in this example is a redundant 1Gbps network link. Under ideal circumstances this would allow us to transfer 1TB of data within approximately 19 minutes at full bandwidth, or about 1 hour when using 30% of the total bandwidth.

After weighing the bandwidth capacity vs. the total amount of data that must be transferred, a migration across the network seems to be the most sensible out of the 3 cold data migration options mentioned above. So, with this information, we now decide upon our exit strategy for each of the services.

M365 365 Service	Business Function	Migration Approach
Exchange Online	Personal mailbox, calendar & contacts	All mailboxes and data are migrated across the network to Exchange on-premise
	Shared mailboxes	
	Email archives	Same as above, but this migration uses a longer time horizon
	Mail flow security EOP, ATP	Continue to use Exchange Online Protection (EOP) in conjunction with Exchange on-premise at the start to facilitate a speedy exit process. In a later stage, start using built-in antimalware and antispam functions of Exchange on-premise or deploy third-party solutions.
MX in/out mail services		
Azure AD (AAD)	AAD User Provisioning	No longer needed, service is discontinued
	Federated authentication	No longer needed, service is discontinued
SharePoint Online	Intranet publishing (websites)	Migration of site collections & document libraries to SharePoint on-premise
	Document collaboration	
OneDrive	Local file storage & sync	OneDrive is redirected to SP on-premise. This allows internal document collaboration.
	Personal document collaboration	
Azure Information Protection (AIP)	Document-level encryption of documents classified as being secret	Keep functionality on-premise using AD-RMS.
Teams	Online collaboration	Stop usage of the pilot, embed some of the use cases within SharePoint.

Step 4: describe future state

General Considerations

As a future state may hold several options, and only one is typically withheld in the exit plans future state description, we must document which assumptions have been taken.

Below we list some assumptions that may apply also to your exit plan:

- The Microsoft 365 service is available when the exit plan is triggered (without the service being available, the exit cannot be executed).
- The current environment is federated, with identities being provisioned and synchronized from an on-premise AD forest to Azure AD (AAD). There is no master data stored within AAD that must be rolled back to on-premise.
- No application integration took place with the service. If this is the case, these applications must be inventoried & plans to redirect these on-premise must be included.

Technology decisions

List up the technology principles and decisions, as well as required hardware and licenses.

Items to consider may include:

- What are the future licensing requirements?

- Compatibility of new service with business requirements: will the new on-premise deployment allow the same functionalities vs. the cloud? Gaps must be analyzed between the cloud service vs. the future on-premise state (see individual service comparison link under Annex 1: Other resources)
- Archive migration: the legal immutable status must be maintained before, during and after the migration
- Encrypted emails: some cloud encryption technologies will no longer exist on-premise; this can break existing functionality or make emails no longer accessible
- Exchange/Skype/Teams federation may pose some challenges and require changes during the migration
- Partner and guest access in SP/OD: external sharing capabilities may need to be foreseen on-premise also, and access issues may arise after the migration to previously shared content
- Remote access requirements: M365 is accessible from anywhere, anytime as a cloud service. However, an alternative solution may introduce some limitations
- Anti-spam, anti-spoofing and anti-malware for inbound/outbound email flows: these flows will likely be affected in an exit scenario and the impact of changes must be assessed

List up the required hard- and software, as in our example below:

Online Service	Infrastructure	Software Solution
Exchange Online	<ul style="list-style-type: none"> ➤ 12 multirole virtual servers (6 in each datacenter) ➤ 24CPU core + 196GB RAM per server ➤ ~80TB of disk storage (40TB in each datacenter without mirror) 	Exchange 2019 running on Windows Server 2019 Datacenter
Exchange Online Protection	No specific hardware is needed for this, mail routing would continue to rely upon EOP online services.	N/A
AAD	Continue to use existing AD Forest (AAD holds a copy of the on-premise AD data)	
SharePoint Online + OneDrive	<ul style="list-style-type: none"> ➤ 12 multirole virtual servers (6 in each datacenter) ➤ 24CPU core + 196GB RAM per server ~40TB of disk storage (20TB in each datacenter without mirror) 	SharePoint 2019 running on Windows Server 2019 Datacenter

People & Support Requirements

High-level description of the supporting environment for the future state.

Team/function	Current roles	New roles
Digital Workplace Help Desk	5 FTE assigned to digital workplace 2 FTE estimated work on M365 related topics as 1 st line workers	Same
M365 Infrastructure Support Team	3 FTE (tenant, EXO, SP/OD, Teams, identity) manage environment and provide 24/7 support in 2 nd line	1 FTE infrastructure support (VM, storage, OS, network...) 4 FTE for 24/7 application support (Exchange & SP/OD only) 2 FTE to work out new (lost) use cases on-premise (alternative for stopped Teams project)
SP Site Administration	2 FTE (site collections)	Same

Team/function	Current roles	New roles
Security Teams	2 FTE (M365 security config, ATP, DLP, training, awareness...)	Same (continue to use EOP)
Compliance	1 FTE (data governance, litigation hold, AIP, eDiscovery, investigations)	1.5 FTE due to the higher skills/time needed for investigation/correlation cross-system

Step 5: high-level migration plan

Going back to our earlier example, we get the following result:

Project Milestones	Project Steps	Owner	Effort [days]	Duration [weeks]
Preparation	Create project plan	...	5	6
	Inventory/update exit plan assumptions (data usage, solution)		5	
	Deep analysis of service usage		5	
	Verify prerequisites and confirm resources for hybrid Exchange		3	
	Verify prerequisites and confirm resources for on-premise SP/OD		3	
	Engage with HR for upstaffing & role transitions		5	
	Acquire partnership/consultancy resources		5	
	Setup project structure		5	
Planning	Create Exchange on-premise deployment plan	...	25	8
	Create SP/OD on-premise deployment plan		20	
	Create data migration schedule		5	
	Create user migration schedule		5	
	Build Exchange on-premise test & acceptance environments		10	
	Build SP/OD on-premise test & acceptance environments		10	
	Prepare network requirements		10	
	Test move mailbox migration for emails		5	
	Test data migration for SharePoint sites		5	
	Prepare deployment of changes to endpoints		5	
	Prepare communication plan for end users		3	
	Create detailed migration plan		5	
	Create plan for permission and other cloud object conversion		10	
	Plan for mail flow change		5	
	Plan co-existence / hybrid environment		5	
	Train operations & support staff		30	
	Train security, compliance, HR admin users		15	
Plan modifications to user objects, authorizations, AD...		20		

	Project management for planning phase		25	
Build solution	Build Exchange on-premise production environment	...	10	4
	Build SP/OD on-premise production environment		10	
	Configure mail flow changes		5	
	Setup hybrid Exchange configuration		3	
	Implement changes needed in authorization systems		5	
	Testing and validation in against production environment		8	
	Project management for build phase		5	
Network upgrades	Execute network changes/upgrades	...	10	3
	Test and validate E2E connectivity, capacity testing		2	
Pilot migration	Migrate limited number of user mailboxes (move mailbox)	...	5	3
	Push changes to client & authorization systems		2	
	Migrate limited size required SP/OD data		3	
	Verify and support co-existence		2	
	Validate migration results & solve open issues		10	
Velocity Migration	Prepare for mass migration	...	3	6
	SP/OD Full data migration (1 weekend)		3	
	Client and authorization updates for OD		3	
	Exchange migration across multiple weekends		15	
	Client and authorizations updates for Exchange		3	
	Help desk & support coordination		10	
Cold Data Migration	Migration of Exchange archives	...	2	2
	Testing and validation		3	
	Project management for migration (all 3 steps)		10	
Wrap-up	Soft delete migrated SPO/EXO data	...	1	2
	Remove unnecessary licenses		1	
	Remove unneeded configuration elements (ADFS proxies...)		3	
	Cancel subscription		1	
Totals			392	34

Disclaimer: these steps and figures are entirely fictive & must not be considered as indicative for real-life exit plans.

In this example the entire exit would take roughly 34 weeks to complete. The estimate of man-days can be translated in both FTE requirements during this period as well as in a hard cost figure.

Step 6: exit plan testing

Please refer to the general section step 6.

Step 7: plan update

Please refer to the general section step 7.

Azure Exit Planning Considerations

Step 1: planning and analysis

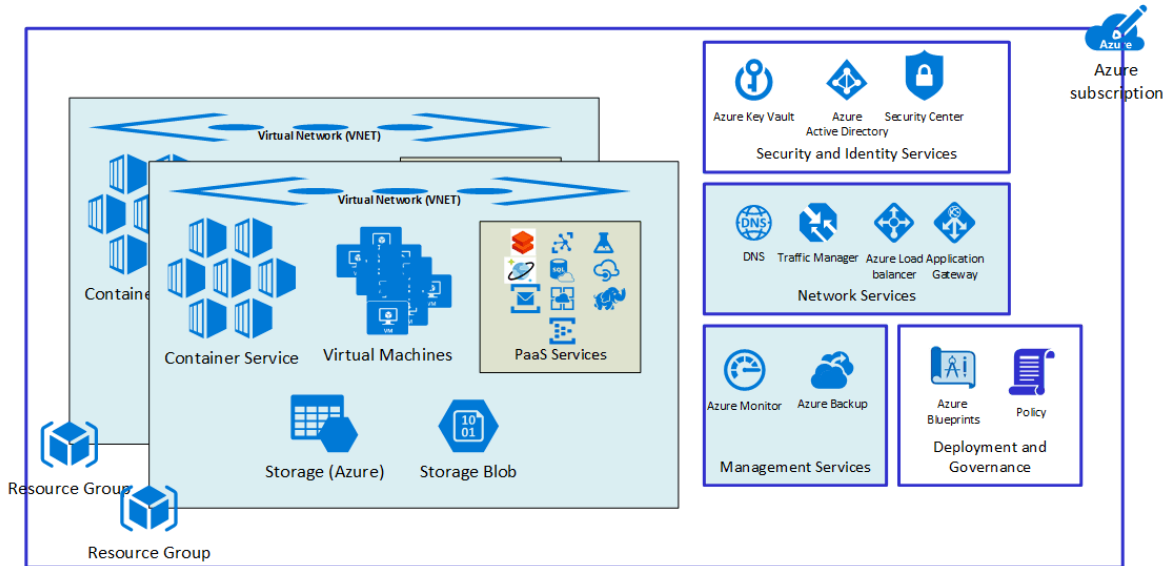


Figure 5 Sample view of Azure Environment

The first step in analysis of the current state environment within Azure is to obtain a full inventory of all resources on a subscription by subscription basis. The data necessary for this inventory can be gathered by [configuring dynamic inventory gathering](#).

A sample list of information for Azure Infrastructure-as-a-Service (IaaS) Virtual Machines (VM) is provided below. It is important to indicate the date at which this information has been gathered so that during later planning phases correct assumptions are made on the amount of data in use.

Note: this list is only a sample and does not contain all the information that a given enterprise would need to gather to fully evaluate their IaaS environment. This includes Business purpose and ownership to enable proper categorization and migration grouping. The enterprise should feel free to add any information they deem necessary in order to properly categorize.

- Information Gathering Date: [July 3, 2019](#)
- Subscription Name: [CustTestSubscription](#)
- Subscription ID: [c506-d97d-49d1-b318-5e71](#)

- For Each IaaS VM:
 - Application: [XYZ-Trad-App-SQL](#)
 - LOB: [Front Office](#)
 - Business Purpose: [SQL Database to support online trading](#)
 - VM Name: [TestInventoryVM001Linux](#)
 - Resource Group: [TestINV-RG](#)

- Location: [East US](#)
- State: [Running](#)
- Operating System: [Linux \(ubuntu 18.04\)](#)
- Size/Type: [Standard D2s v3 \(2 vcpus, 8 GiB memory\)](#)
- Public IP Address: -
- DNS Name: -
- Private IP Address: [10.0.0.4](#)
- VNET/Subnet: [TestInv-RG-vnet/default](#)
- Subnet Mask: [10.0.0.0/24](#)
- OS Disk:
 - Size: [30 GiB](#)
 - Storage Account Type: [Standard SSD](#)
 - Encryption: [Enabled](#)
 - Host Caching: [Read/Write](#)
- Data Disk 1:
 - Size: [2048 GiB](#)
 - Storage Account Type: [Premium SSD](#)
 - Encryption: [Enabled](#)
 - Host Caching: [Read-only](#)
- Availability Set: [TestInv-AVSET](#)
- Availability Zone: -

- For Container Instances
 - Container Name: [testinvcontainer001](#)
 - Resource Group: [TestINV-RG](#)
 - Location: [East US](#)
 - Image Type: [Public](#)
 - Image Name: [testimagereg.azurecr.io/hello-world](#)
 - OS Type: [Linux](#)
 - Number of CPU Cores: [1](#)
 - Memory: [1.5](#)
 - Restart Policy: [On Failure](#)

- Container Registry
 - Resource Group: [TestINV-RG](#)
 - Location: [East US](#)
 - Login server: [testimagereg.azurecr.io](#)

- For Each PaaS Service (Example of Azure SQL)
 - SQL Database Name: [TestInventoryAzureSQL001](#)
 - Resource Group: [TestINV-RG](#)
 - Location: [East US](#)
 - State: [Running](#)
 - Server Name: [testinventoryserversql001.database.windows.net](#)
 - Service Specific Configurations
 - Elastic pool: [Show Here](#)
 - Connection Strings: [Show Here](#)

- Geo-Replication: [Show Here](#)

- For Each PaaS Service (Example of Azure Cosmos DB)
 - Resource Group: [TestINV-RG](#)
 - Read Location: [East US](#)
 - Write Location: [East US](#)
 - URI: <https://testinvcosmos.documents.azure.com:443/>
 - API: [SQL](#)
 - State: [Online](#)

Once a solid inventory of the components and services is completed the resources can be grouped by: Business function, Line of Business, Deployment Type (IaaS, Container, PaaS), Location, Function Etc. The grouping methodology can take any form the enterprise finds useful. The above example provides enough information to group as necessary. Additionally, all Management, Security, Networking, Identity, and other platform services must be assessed so that these functions can be migrated to the target environment.

[Subscription](#) structure, [management group](#) layout and effective [policies](#) should be assessed in order to determine the controls being used and required for set up in the target environment.

Use [Azure Resource graph](#) to quickly and efficiently query across Azure subscriptions. Analyze your cloud inventory using complex querying launched programmatically or from the Azure portal. And assess the impact of applying policy in extensive cloud environments.

Step 2: risk assessment

As discussed in the introduction of this paper (myth 1), exit planning involves the ultimate risk mitigation but should never be your primary measure in dealing with CSP service failures.

It is much more important to choose the most appropriate Azure architecture blueprint for your business deployment. By making the right design choices upfront, you will increase overall resiliency of your solution and inherently reduce risk of failures.

Our SLA's for Azure are here: <https://azure.microsoft.com/en-us/support/legal/sla/>

Azure Blueprints are here: <https://docs.microsoft.com/en-us/azure/governance/blueprints/overview> and hold some specific security and compliance design blueprints for the financial services industry aligned with industry standards such as FFIEC and PCI DSS.

These blueprints are important because the SLA you receive for your business service fully depends upon the chosen deployment model. For instance, at time of writing of this whitepaper the [SLA for IaaS VM's](#) could go up to 99.99% depending upon which deployment model you pick:

- For all Virtual Machines that have two or more instances deployed across two or more Availability Zones in the same Azure region, we guarantee you will have Virtual Machine Connectivity to at least one instance at least 99.99% of the time.
- For all Virtual Machines that have two or more instances deployed in the same Availability Set, we guarantee you will have Virtual Machine Connectivity to at least one instance at least 99.95% of the time.
- For any Single Instance Virtual Machine using premium storage for all Operating System Disks and Data Disks, we guarantee you will have Virtual Machine Connectivity of at least 99.9%.

Going back to the example above for a database system like CosmosDB, and considering the possible deployment models, you can even achieve a 99.999% [SLA for a CosmosDB](#) dependent business service:

- Azure Cosmos DB is Microsoft’s globally distributed multi-model database service. It offers turnkey global distribution across any number of Azure regions by transparently scaling and replicating your data wherever your users are. The service offers comprehensive 99.99% SLAs which covers the guarantees for throughput, consistency, availability and latency for the Cosmos DB Database Accounts scoped to a single Azure region configured with any of the five Consistency Levels or Database Accounts spanning multiple Azure regions, configured with any of the four relaxed Consistency Levels. Azure Cosmos DB allows configuring multiple Azure regions as writable endpoints for a Database Account. In this configuration, Cosmos DB offers 99.999% SLA for both read and write availability.

Coming back to the exit plan, when assessing triggers that may lead to an exit, most often repeated SLA breaches are one of the possible events that would drive a customer to invoke an exit plan. This is explained in the general description of Step 3: decide exit strategy.

We encourage customers to define formal exit plan triggers associated with repeated SLA failures. While the exit plan will not provide help at the time of an individual failure, it will provide a structural solution in case of repeated failures leading to a loss of trust in the CSP.

Step 3: decide exit strategy

Each of the Azure deployment above, as well as others relevant to your business, will require a different means of exit from the Azure cloud. Although no one solution can cover all the use cases within a given deployment the following table will provide a basic set of guidelines for most cases.

In section 5, the steps given are only for the extraction of usable data and configurations from Azure to an intermediate storage point, not the installation into another platform.

Deployment Type	Exit Strategy
IaaS VM	Backup and Recovery outside of the Azure Cloud
Containerized workload	Container and Cloud Service Migration
PaaS Service	Create service on the target platform and migrate data

Step 4: describe future state

Three possibilities exist for the future state upon exit from a given public cloud environment:

- 1. Migration to another cloud**
 - a. Migration to another public cloud environment
 - b. Migration to a private cloud environment

The key to migration between public or private cloud environments is the matching of components or services from the origin cloud to the target cloud. For each component or service a corresponding component or service must be deployed, configured and the data migrated to the target environment. For example, a SQL database will run in almost any environment with minor configuration changes.

- 2. Migration to an on-premise virtualized or non-virtualized environment**

The difference between migration to another cloud and migration to on-premise is the lack of a services layer. The services layer includes management, monitoring, security, managed storage, and networking. Each of these services and others need to be integrated with on-premise processes, procedures and technology.

For any of the above scenarios a full architecture, design, implementation plans and testing plans, need to be developed to enable the efficient migration of components and services to the target. While this is beyond the scope of an exit planning document, the necessary steps should be highlighted into this document so that the impact can be assessed in terms of people, infrastructure, skills and migration time requirements.

Step 5: high-level migration plan

As deployment methodologies will vary greatly depending on the target environment, this section will focus on the migration of Virtual Machine, Services and Data out of the Azure cloud to be accessed by the target. [Azure Backup](#) and [Azure Snapshot](#) technologies are intended for retention and protection of data within the Azure Cloud. In the cases below that backup and recovery to another environment are needed, it is suggested that a third-party backup/recovery toolset that supports Azure and the target environment be used (example [Veritas Netbackup for Azure](#)). The estimated times are per instance and are based on the following example environment:

- 50 GB SQL DB
- 1 Gbps ExpressRoute

No estimation is made for the target environment setup but the following, and much more, must be in place before the workload can be moved:

- Networking
- Storage
- Workload Infrastructure (Physical, Virtualized, or Container)
- Security

Deployment Type	Estimated Timeframe	High Level Exit Plan Steps
IaaS VM	Backup = 15 Mins Data Xfer = 8 Mins	<ul style="list-style-type: none"> • For each VM to be migrated <ol style="list-style-type: none"> a. Backup all data, Operating System and Configurations b. Store the Backed-up Data in a target environment accessible location. • On the target environment <ol style="list-style-type: none"> a. Unload the data b. Re-install the OS and applications in the target environment
Containerized workload	Backup = 15 Mins Data Xfer = 8 Mins	<ul style="list-style-type: none"> • For Each container (Azure Container Services, Azure Kubernetes Services) <ol style="list-style-type: none"> a. Backup all data • On the target environment <ol style="list-style-type: none"> a. Ensure that the container host OS and API set are identical on both the origin and the target environment.

Deployment Type	Estimated Timeframe	High Level Exit Plan Steps
		<ul style="list-style-type: none"> b. Register the workload with the target container service c. Redeploy the container to the target
PaaS Service	Backup = 15 Mins Data Xfer = 8 Mins	<ul style="list-style-type: none"> • For each PaaS service (example – Azure SQL Database) <ul style="list-style-type: none"> a. Re-create service and management layer in target platform b. Choose and deploy a corresponding service or software environment in the target environment • For each use of the PaaS service <ul style="list-style-type: none"> a. Backup all data b. Store the data in a target environment accessible location

Figure 6 Azure Exit Plan High-Level Steps

Step 6: exit plan testing

Please refer to the general section step 6.

Step 7: plan update

Please refer to the general section step 7.

Dynamics 365 Customer Engagement Considerations

There is no written-out guidance provided walking across all 7 steps as in the M365 example above. Instead we focus on helping you to decide the most appropriate exit strategy under step 3.

Step 3: decide exit strategy

As Microsoft's online Dynamics 365 offerings advance more quickly than on-premises offerings, it is not possible to directly migrate from online to on-premises environments (lift-and-shift style operations). There are 3 alternative ways to establish this.

Option 1: Bulk data migration via data lake

You could export the data to data lake as an intermediate step, then import the data in the lake into other systems. The process for exporting to Azure Data Lake Storage is documented [here](#).

Option 2: Event-driven data movement using Microsoft Power Automate

Power Automate is a service that helps you create automated workflows between your favorite apps and services to synchronize files, get notifications, collect data and more. See: <https://docs.microsoft.com/en-us/power-automate/>. As data is created, updated, or deleted in your environment, you can trigger flows to copy data, or invoke similar actions, against any of the 300+ connectors available.

Option 3: Custom data movement via oData

All Dynamics 365 Customer Engagement data is accessible via oData APIs. You can write custom code which reads data from the oData APIs and writes the data into the destination system of your choice. API documentation is available [here](#).

Annex 1: Other Resources

General Information	
Product Licensing Terms	General termination conditions for Microsoft 365 are documented in the Online Services Terms (OST). In addition, specific conditions may apply for certain industries such as the Financial Services Industry. Please check your contract accordingly for the applicable licensing terms.
Microsoft 365 Service Descriptions	Service descriptions of all Microsoft 365 services, including 1-by-1 feature comparison against on-premise alternatives where available.
MS Exchange Service Description	
SharePoint Service Description	
Skype for Business Service Description	
Office Online Server documentation	Office Online Server is a product that allows browser-based usage of office documents in your own data center, similar as this exists in Microsoft 365. You may want to deploy this as part of the rollback.
Online Services SLA	The official SLA documentation for all online services.
Azure SLA	Service Level Agreement with additional details for Azure IaaS/PaaS Online Services

Technical Migration Scenario support	
Azure Information Protection (AIP) exit planning guidance	A blogpost describing "How to prepare an Azure Information Protection "Cloud Exit" plan"
SharePoint Hybrid	Starting page to explore different hybrid scenarios for SharePoint services, including OneDrive
Deploying hybrid Exchange and how to move mailboxes	Documentation on how to deploy a hybrid MS Exchange environment which would offer the most flexible way to onboard M365 as well as to facilitate an exit if necessary.
Azure AD Connect Attributes Synchronized	List of attributes relevant to M365 Services that may need to sync across on-premise and cloud services during a cloud exit scenario when running in hybrid setup.
Skype Hybrid Planning Guide	Documents the scenarios for running Skype in hybrid mode
M365 Resource Limits & Service Throttling	As part of our efforts to isolate tenants from each other we apply limits to the amount of resources that can be consumed at any specific time. These limits may present barriers for large data transfers, e.g. in case of a service exit scenario.

External resources	
EBF Guidance on Exit Plan Testing	The European Banking Federation (EBF) published guidance on the testing of exit plans in 2020, which may be a helpful underlying resource to step 6 of this document.
SWIPO	SWIPO (Switching Cloud Providers and Porting Data), is a multi-stakeholder group facilitated by the European Commission, in order to develop voluntary Codes of Conduct for the proper application of the EU Free Flow of Non-Personal Data Regulation / Article 6 "Porting of Data"

Data Transfer Project	The Data Transfer Project was formed in 2017 to create an open-source, service-to-service data portability platform so that all individuals across the web could easily move their data between online service providers whenever they want. This project is still under development at the time this whitepaper was published but may offer valuable resources in support of creating an exit strategy over time.
Exoprise	Third-party that provides monitoring services on cloud providers such as Microsoft.
Draft standard ISO/IEC 19941	When planning data migration strategies, ISO standard 19941 highlights some of the complexities, challenges and facets that must be considered around data portability.
Draft standard ISO/IEC 19944	This draft data taxonomy talks about data flows, categories & use

Azure Tools	
Azure Backup	Azure Backup service
Azure Snapshot	Azure VM Snapshot Service
Azure Site Recovery	Azure VM Migration and DR tool set
Azure Container Services , Azure Kubernetes Services	Azure container service options
Example Azure PaaS Services	Azure SQL Database Azure Cosmos Database HDInsight Azure Databricks Azure ML Service Event Hubs

Annex 2: Regulation Overview

Australian Prudential Regulation Authority (ARPA)

The [Australian Prudential Standard CPS 231](#) of July 2017 requires regulated financial organizations, as part of their assessments in case of material outsourcing activities, to ‘*have developed contingency plans that would enable the outsourced business activity to be provided by an alternative service provider or brought in-house if required*’ (article 26).

During the assessment of the options for outsourcing it must also take into account ‘*contingency issues in accordance with [Prudential Standard CPS 232 Business Continuity Management \(CPS 232\)](#) should the outsourced activity need to be brought in-house*’ (article 27).

European Banking Authority (EBA)

On 25 February 2019, the [EBA published its final guidelines on outsourcing arrangements](#) which includes a requirement for financial institutions to develop exit strategies when outsourcing critical or important functions.

In these EBA guidelines under Title IV Section 14, the following paragraphs are relevant for exit planning if your financial institution is present in the European Union:

106. ***Institutions and payment institutions should have a documented exit strategy when outsourcing critical or important functions that is in line with their outsourcing policy and business continuity plans, taking into account at least the possibility of:***
 - a. *the termination of outsourcing arrangements;*
 - b. *the failure of the service provider;*
 - c. *the deterioration of the quality of the function provided and actual or potential business disruptions caused by the inappropriate or failed provision of the function;*
 - d. *material risks arising for the appropriate and continuous application of the function.*

107. ***Institutions and payment institutions should ensure that they are able to exit outsourcing arrangements without undue disruption to their business activities, without limiting their compliance with regulatory requirements and without any detriment to the continuity and quality of its provision of services to clients. To achieve this, they should:***
 - a. *develop and implement exit plans that are comprehensive, documented and, where appropriate, sufficiently tested (e.g. by carrying out an analysis of the potential costs, impacts, resources and timing implications of transferring an outsourced service to an alternative provider); and*
 - b. *identify alternative solutions and develop transition plans to enable the institution or payment institution to remove outsourced functions and data from the service provider and transfer them to alternative providers or back to the institution or payment institution or to take other measures that ensure the continuous provision of the critical or important function or business activity in a controlled and sufficiently tested manner, taking into account the challenges that may arise because of the location of data and taking the necessary measures to ensure business continuity during the transition phase.*

108. ***When developing exit strategies, institutions and payment institutions should:***
 - a. ***define the objectives of the exit strategy;***
 - b. ***perform a business impact analysis that is commensurate with the risk of the outsourced processes, services or activities, with the aim of identifying what human and financial resources would be required to implement the exit plan and how much time it would take;***

- c. **assign roles, responsibilities and sufficient resources** to manage exit plans and the transition of activities;
- d. **define success criteria** for the transition of outsourced functions and data; and
- e. **define the indicators to be used for the monitoring of the outsourcing arrangement** (as outlined under Section 14), including indicators based on unacceptable service levels that should trigger the exit.

European Insurance and Occupational Pension Authority (EIOPA)

On 6 February 2020, EIOPA also issued [Guidelines on outsourcing to cloud service providers](#) which include specific requirements on termination rights and exit strategies.

Guideline 15 – Termination rights and exit strategies

55. *In case of cloud outsourcing of critical or important operational functions or activities, within the cloud outsourcing agreement the undertaking should have a clearly defined exit strategy clause ensuring that it is able to terminate the arrangement, where necessary. The termination should be made possible without detriment to the continuity and quality of its provision of services to policyholders. To achieve this, the undertaking should:*

- a. *develop exit plans that are comprehensive, service based, documented and sufficiently tested (for example, by carrying out an analysis of the potential costs, impacts, resources and timing implications of the various potential exit options);*
- b. *identify alternative solutions and develop appropriate and feasible transition plans to enable the undertaking to remove and transfer existing activities and data from the cloud service provider to alternative service providers or back to the undertaking. These solutions should be defined with regard to the challenges that may arise because of the location of data, taking the necessary measures to ensure business continuity during the transition phase;*
- c. *ensure that the cloud service provider adequately supports the undertaking when transferring the outsourced data, systems or applications to another service provider or directly to the undertaking;*
- d. *agree with the cloud service provider that once retransferred to the undertaking, its data will be completely and securely deleted by the cloud service provider in all regions.*

56. *When developing exit strategies, the undertaking should consider the following:*

- a. *define objectives of the exit strategy;*
- b. *define the trigger events (for example, key risk indicators reporting an unacceptable level of service) that could activate the exit strategy;*
- c. *perform a business impact analysis commensurate to the activities outsourced to identify what human and other resources would be required to implement the exit plan and how much time it would take;*
- d. *assign roles and responsibilities to manage exit plans and transition activities;*
- e. *define success criteria of the transition.*

European Securities and Markets Authority (ESMA)

On 3 June 2020 ESMA as a third European supervisory authority opened a [public consultation on outsourcing to cloud service providers](#). At the time of writing of this whitepaper this document is still in consultation format, but it provides insight into the expectations set forward by this authority. Guideline 5 of this principles based document covers exit strategies:

44. *In case of outsourcing of critical or important functions, a firm should ensure that it is able to exit cloud outsourcing arrangements without undue disruption to its business activities and services to its clients, and without any detriment to its compliance with the applicable legal requirements, as well as the confidentiality, integrity and availability of its data. To achieve this, a firm should:*

- a) *develop and implement exit plans that are comprehensive, documented and sufficiently tested. These plans should be updated as needed, including in case of changes in the outsourced function;*
- b) *identify alternative solutions and develop transition plans to remove the outsourced function and data from the CSP and, where applicable, any sub-outsourcer, and transfer them to the alternative CSP indicated by the firm or directly back to the firm. These solutions should be defined with regard to the challenges that may arise from the location of the data, taking the necessary measures to ensure business continuity during the migration phase;*
- c) *ensure that the cloud outsourcing written agreement includes an obligation for the CSP to orderly transfer the outsourced function and all the related data from the CSP and any sub-outsourcer to another CSP indicated by the firm or directly to the firm in case the firm activates the exit strategy;*
- d) *ensure that any data removed or transferred is securely deleted from the systems of the CSP and, where applicable, of any sub-outsourcer (for example, by requesting a written confirmation by the CSP).*

45. *When developing the exit plans and solutions referred to in points (a) and (b) above ('exit strategy'), the firm should consider the following:*

- a) *define the objectives of the exit strategy;*
- b) *define the trigger events that could activate the exit strategy. These should include at least the termination of the cloud outsourcing arrangement at the initiative of the firm or the CSP and the failure or other serious discontinuation of the business activity of the CSP;*
- c) *perform a business impact analysis that is commensurate to the function outsourced to identify what human and other resources would be required to implement the exit strategy;*
- d) *assign roles and responsibilities to manage the exit strategy;*
- e) *test the exit strategy, using a risk-based approach;*
- f) *define success criteria of the transition.*

46. *The firm should include indicators of the trigger events of the exit strategy in its ongoing monitoring and oversight of the services provided by the CSP.*

Hong Kong Monetary Authority Supervisory (MAS)

Supervisory Policy Manual SA-2 of December 2001 requires authorized institutions to maintain and test contingency plans. In establishing contingency plans, authorized institutions is required to 'consider' availability of alternate services or possibility of moving back to on-premise deployments.

The descriptions in this document is aligned with SA-2 requirement and help institutions to establish and maintain the contingency plans compliant with SA-2.

Japan Financial Industry Information Systems (FISC)

The Center for Financial Industry Information Systems (FISC) was established in November 1984 as an incorporated foundation under the approval of the then Minister of Finance. In collaboration with its member institutions, the Financial Services Agency and the Bank of Japan, FISC has established several guidelines for the promotion of security measures on financial institutions information systems. These guidelines have been voluntarily observed by most financial institutions in Japan.

For the context of this document, the *Manual for the Development of Contingency Plans* is relevant and key steps have been incorporated into this document.

Monetary Authority of Singapore (MAS)

The [Guidelines on Outsourcing](#) of July 2016 (revised October 2018) require a financial institution to consider 'worst case scenarios' as part of its business continuity plans. Also, the guidelines explicitly require business continuity management plan and exit strategy when a financial institution has outsourcing arrangement outside Singapore. nstitution has outsourcing arrangement outside Singapore.

Swiss Bankers Association (SBA)

In March 2019 the Swiss Bankers Association (SBA) issued [Legal and regulatory guidelines for the use of cloud services](#) by banks and securities dealers in the context of FINMA-regulated outsourcing, and these also include Measures to secure the availability and return of information on page 36:

53. The institution should be able to access any protected information that is stored or processed abroad or in Switzerland at any time from Switzerland. The provider should undertake to continue to deliver the cloud services to the institution, a successor company or rescue company and, where applicable, FINMA if the institution is in recovery or resolution, to the extent that such access from Switzerland to information abroad or in Switzerland is assured as a result.
54. The provider should undertake to return the protected information to the institution, a successor company or rescue company at any time as part of assistance with termination, if the institution is in recovery or resolution and on the instructions of the institution or FINMA, provided the provider has the means¹² and knowledge¹³ to do so. In this case, the provider should transfer the protected information back in a standardised, machine-readable format.
55. If the provider uses proprietary solutions that result in lock-in effects, the provider should declare its willingness to support the institution with a migration to other solutions or with licensing such solutions.

UK Financial Conduct Authority (UK FCA)

A firm should:

- a) have exit plans and termination arrangements that are understood, documented and fully tested;
- b) know how it would transition to an alternative service provider and maintain business continuity;
- c) have a specific obligation put on the outsourcing provider to cooperate fully with both the firm and any new outsource provider(s) to ensure there is a smooth transition;
- d) know how it would remove data from the service provider's systems on exit;
- e) monitor concentration risk and consider what action it would take if the outsource provider failed.

UK Bank of England's Prudential Regulatory Authority (UK PRA)

On 5 December 2019 the UK PRA released [Consultation Paper 30/19 on Outsourcing and third party risk management](#) extensively covering exit planning under section 10 Business continuity and exit plans. This document holds multiple requirements and makes a distinction between stressed and non-stressed exits.

We copied the most important excerpts below:

10.1 For each material outsourcing arrangement, the PRA expects firms to develop, maintain and test a:

- *business continuity plan; and*

¹² Such as encryption keys.

¹³ Especially where cloud services as part of IaaS or PaaS are concerned, the provider may have no knowledge of the architecture chosen by the institution and / or the components used by the institution.

- *documented exit strategy, which should cover and differentiate between situations where a firm exits an outsourcing agreement:*
 - *in stressed circumstances, eg following the failure or insolvency of the service provider (stressed exit); and*
 - *through a planned and managed exit due to commercial, performance or strategic reasons (non-stressed exit).*

10.2 The PRA's primary focus when it comes to business continuity plans and exit strategies is on the ability of firms to deliver important business services provided or supported by third parties in line with their impact tolerances in the event of disruption. Consequently, notwithstanding the importance of effectively planning for non-stressed exits, the main focus of this chapter is on business continuity and stressed exits.

Stressed exits

10.8 Firms' exit plans should cover stressed exits and be appropriately documented and tested as far as possible.

10.9 A key objective of the stressed exit part of exit plans is to provide a last resort risk mitigation strategy in the event of disruption that cannot be managed through other business continuity measures, including those mentioned in the previous section, eg the insolvency or liquidation of a service provider.⁴⁷

10.10 The PRA does not prescribe or have a preferred form of exit in stressed scenarios. Its focus is on the outcome of said exit, ie the continued provision by the firm of important business services provided or supported by third parties, rather than the method by which it is achieved.

10.11 The PRA does, however, expect firms to consider all potentially viable forms of exit in a stressed exit scenario, which may include:

- *bringing the data, function or service back in-house/on-premise;*
- *transferring the data, function or service to an alternative or back-up service provider; or*
- *any other viable methods.*

10.12 The PRA expects firms to give meaningful consideration to all available tools that can facilitate an orderly stressed exit from a material outsourcing arrangement. These tools are constantly evolving, in particular in technology outsourcing, including Cloud, and may include:

- *new potential service providers;*
- *technology solutions and tools to facilitate the switching and portability of data and applications; and*
- *industry codes and standards.*

10.13 Firms should also actively consider temporary measures that can help ensure the ongoing provision of important business services following a disruption and/or a stressed exit, even if these are not suitable long-term solutions, eg contractual arrangements allowing for continued use of a service or technology for a transitional period following termination.

Governance of business continuity plans and exit plans

10.14 Firms should develop their business continuity and exit plans, in particular for stressed exits, during the pre-outsourcing phase once they have determined that a planned outsourcing arrangement is material (see Chapter 5). Doing so will enable them to:

- *use the due diligence process to identify potential alternative or back-up service providers;*

- estimate the cost, resourcing and timing implications of the proposed business continuity or exit plan in both stressed and non-stressed scenarios as part of the risk assessment;
- identify data they may need to access, recover or transfer as a priority in a disruption or stressed exit; and
- define the key KPIs and key risk indicators (KRIs) which, if breached, may trigger an exit (both stressed and non-stressed).

10.15 Firms should assign clear roles and responsibilities for business continuity and exit plans. Subject to proportionality, they may establish cross-disciplinary teams to develop, document, test and execute their business continuity and exit plans, especially in stressed scenarios (which may include communicating with the PRA and other relevant stakeholders in the event of disruption). Based on the size and complexity of the firm, these teams may include relevant business lines, control functions, technical experts (eg IT specialists) and be chaired by an SMF. Firms should also allocate responsibility for signing off business continuity and exit plans, including updates thereafter, and the decision to activate them.

10.16 When developing business continuity and exit plans, firms should define the objectives of the plan, including what would constitute successful business continuity or a successful exit in both stressed and non-stressed scenarios, by reference to measurable criteria such as costs, functionality, time and the firm's impact tolerances for important business services.

10.17 Firms should take reasonable steps to test exit plans; in particular, those relating to stressed exits. The extent and nature of testing will vary depending on the type of outsourcing arrangement and corresponding exit plan. For instance, a firm running a hybrid Cloud structure may take into account the potential back-up functions located in its private Cloud elements. Likewise, a firm that keeps backup copies of data which it has outsourced to the Cloud outside the Cloud environment may focus its testing on assessing the ongoing consistency of both sets of data and reconciling them as appropriate. Firms should also assess and take reasonable steps to manage any operational risks which may be caused or increased by the actual testing (eg data theft).

10.18 Business continuity and exit plans should be reviewed periodically to take into account developments that may change the feasibility of the business continuity measures or an exit, eg:

- an increase in the number of availability zones or regions offered by a current service provider;
- changes to the firm's business requirements;
- the emergence of new, potentially viable alternative providers; and/or
- developments in technology or other tools to facilitate the porting of data and applications, eg among Cloud providers or between firms' on-premise environments and the Cloud.