Documentation for data centre migrations

Data centre migrations are part of the normal life cycle of a typical enterprise. As organisations expand, many reach a point where maintaining multiple, distributed data centres becomes increasingly complex and expensive to support. Integration of systems, applications and infrastructure, together with consolidation of staff and support services becomes a key driver for improving efficiency, and reducing complexity and costs.

The chances are that you will encounter an enterprise undergoing this process at some stage in your career as a technical communicator. I was recently involved in a complex migration project that involved consolidating 18 data centres, located throughout Europe, into a single, centralised data centre. This article describes some of the strategies and methods we used for collecting and managing the information required for the data centre migration.

Migrations are highly complex and what is described in this article is only intended to provide an introduction to this topic.

Adopting a structured approach to data collection

Reliable information about the current data centre infrastructure is key to facilitating a smooth migration and avoiding downtime to critical business systems. It is often the role of technical communicators to collect this information, which is then used by technical architects and project managers to plan the migration and set up new infrastructure.

For the technical communicators involved, knowing where to start and how to manage the information gathering can be a challenge. A structured approach to your project can help ensure its success. The steps described in this article provide a methodological approach to this task.

Step 1. Set clear project expectations

Set clear expectations about your role at the start of the project. (This is essential, as you may be working with colleagues who have overlapping responsibilities.) One way to do this is by producing a project brief, which is a short document that outlines your role, the scope of the deliverables, and time scales for completion. Have this approved before commencing any work.

Once you have a high-level agreement outlining your role on the project, you can start working towards a detailed project plan, which outlines:

- What data you will be collecting
- How you will collect the data
- When you will need to do this by

There may be dependencies, for example, the availability of subject matter experts who can provide you with the information you need, or work that needs to be done by third parties before you can commence. Your project plan should make provision for any foreseeable dependencies.

Your plans should also provide sufficient flexibility to accommodate changes to requirements, processes and tools during the course of the project.

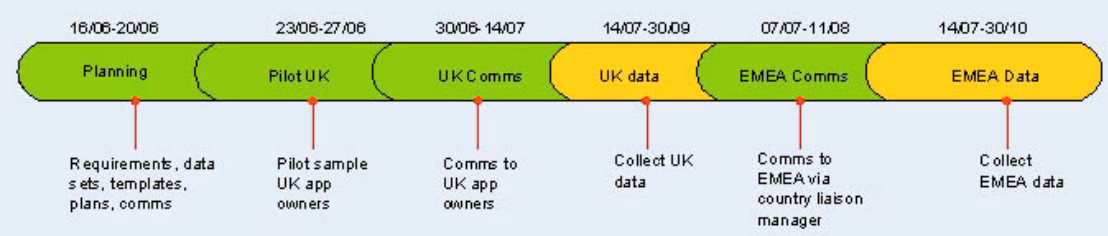

Figure 1. Simplified example of a project plan

Step 2. Clarify what data you need to collect

Ensure that key business stakeholders are engaged and given an opportunity to provide their feedback as to what information they want collected. Key stakeholders should represent IT infrastructure, IT security, business continuity planning, application architecture, service management and any other business areas that are either impacted by the migration or involved in migration planning.

Note: *Failure to engage at the planning stage with relevant business stakeholders could result in your team having to repeat the data gathering exercise at a later stage, as you did not capture all the information required. This could add significantly to project costs and cause delays to migration.*

It is sometimes tricky to determine the level of detail you need when collecting information. Asking for too much detail could significantly slow down migration planning. Ask for too little detail and you may not have enough information to plan the migration.

The following examples illustrate the type of information you may need to collect. The precise details will depend on the organisation and the nature of the project.

Server data

You will need to collect data on the physical infrastructure that is target for migration. Table 1 outlines typical server information:

Table 1. Server details
• Memory, disk and processor specifications
• Operating systems
• Services
• Serial numbers (e.g., for hardware and operating system)
• Host names and IP addresses
• Maintenance agreements
• Age of the machine and any service level agreements (SLAs)
• Physical dimensions — size, weight

Table 1. Server details
• Rack location/blades • Peripherals: monitors, keyboards or other equipment • Heating and ventilation requirements • Security requirements • Power supply requirements • Number of available ports/slots

Physical dimensions, and heating and cooling considerations will determine the available spacing and new room arrangements. Security requirements will determine who has access to the new equipment and the type of access to the room. Old servers that have been running for a number of years may not restart easily if shut down. In this case, migration planners might consider purchasing new kit. If applications on current servers will be installed on new servers or virtualised, the planner needs to ensure that the new servers have sufficient ports. New maintenance agreements may need to be arranged, especially if the server will be running from a different region.

Migration planning, based on the information you collect, will determine whether a server can be lifted and shifted to a new location, or whether a new version needs to be built on new hardware.

Application data

To avoid confusion, it helps to have a clear business definition of what constitutes an 'application'. For example, how does your business see this as differing from a service or process?

Collecting application information is a complex and time-consuming task. There may be numerous applications installed on each physical server. Each application may in turn connect to several other applications. Some of these applications may be hosted on site, others may be hosted on sites belonging to third parties.

Application discovery can be complicated if an enterprise is not aware of all its current applications. A large enterprise can run hundreds of applications in each data centre. In this case, one of your first tasks may be to create an up-to-date list of applications.

Table 2 provides examples of some of the application details you may need.

Table 2. Application details
• Application purpose/function • Service Level Agreements (SLAs) • Business criticality measures • Business services/products and processes supported

Table 2. Application details
• Maintenance windows, change freezes and backup times • Users and customers supported • Application sponsors, subject matter experts, commercial owners and delivery owners • Application architecture and network topology • Upstream and downstream applications connected to this application • Ports and firewalls • Application interfaces (e.g., fax, mail relay, printing) • Hosting arrangements: such as production, disaster recovery, test and development environments, their IP addresses and whether they are physical or virtual • Location (data centre/country where installed) • Licence details – type, serial numbers • Any planned changes/upgrades • Operating systems, databases and versions • Messaging (e.g., MQ, FTP) • Data storage and archiving requirements • Supported batch processes and work flows

Identify a contact person, owner and sponsor for each application. Make sure you have a clear definition of owner, contact and sponsor as these are different roles. The contact provides information. The owner must have decision-making ability, be accountable for the decisions made and specify what testing will be needed. The sponsor represents the business or client. Often they do not know the technical details of the application. All roles are important. It takes time to clarify roles and is not always straightforward.

Prioritise your applications — identify which applications should be targeted first for data collection. You could base this priority on regional or project considerations — the plan might be to migrate each region or data centre in stages. Additionally, you would want to prioritise your information collection to focus first on applications that are key to the business. This may involve discussions with key business stakeholders, who have a good understanding of the core applications in their business area.

You may also want to prioritise the information you collect, to distinguish essential information from nice to have. The key application information you need includes the application owner, the criticality of the application, maintenance windows, SLAs, and how the application links to other applications and to external clients. This information is used to determine the impact and timing of the application migration.

For example, if application A connects to business-critical applications B, C and D, then the migration planners may decide to move all these applications at the same time. Available maintenance windows determine when they can be moved. SLAs determine how much system downtime can be tolerated and who can approve extensions. You will need a business owner to confirm any migration decisions made.

Adopting a staged approach

When a large amount of information needs to be collected, it makes sense to break down the information collection into stages. For example:

Stage 1: Collect high-level information on current infrastructure and applications. For example, what's out there? What does it run on? How important is it to the business? Does it need to be migrated?

Stage 2: Collect detailed information on relevant applications and servers that need to be migrated.

This flexible approach was one we adopted on our project. It enabled a quick initial information gathering exercise to be completed, which then lead to a detailed discovery. This also ensured that subject matter experts were not overwhelmed with requests for information at a single point in time, minimising the impact on regular business activities.

Step 3. Define the methods and tools to be used for data collection

The methods and tools used for collecting information depend on the nature of both the data and the organisation. You can use a combination of the examples shown in Table 3.

Step 4. Determine how information will be stored, maintained and presented

It is essential that the information you collect is stored in a format that is easily updateable and retrievable by relevant areas of the business. If this information can be entered into a database, this makes it easier to generate reports and retrieve the data, as well as providing version control and tracking.

Companies that specialise in migration services provide database tools specifically designed for this purpose.

Step 5. Identify areas in the business that need to be contacted

The information you need to collect may span different business areas and responsibilities.

You will need to discover how the infrastructure is managed, as IT systems will be supported and owned by different departments. For example, mainframe, Windows, Unix systems and telecommunications equipment are usually managed by separate teams. This becomes even more complicated when businesses are spread across countries and regions, as some regions may have wider responsibility for elements of the infrastructure.

Clarify what business areas are involved and obtain nominated people who will be your contact and sponsor in each area.

Step 6. Define a clear communication strategy

Clear stakeholder management is vital in large migration projects. You may need to communicate with key decision-makers and subject matter experts located across different business areas and regions. Providing a consistent message and approach will maximise co-operation with the data collection effort.

A clear communication plan should include:

- The communication process (who is involved, and how and when you will be communicating)
- Response times for returning information requested
- Escalation procedures
- E-mail and communication templates

The key to an effective communication plan is to ensure that it fits into the organisation's existing structure and work practices, and can be easily implemented.

Table 3. Data collection methods and tools	
Electronic data collection	Automated scripts or queries that can be run on machines in the network to collect basic information on operating systems, IP addresses, applications, services, protocols, data volumes and data flows. Software may already be installed that monitors and provides statistics on the network.
Questionnaires	Provide a standard format for collecting information that cannot be obtained electronically. They may be presented as simple Excel spreadsheets, with fields that contacts need to complete, or HTML forms that can be submitted online.
Diagrams	Describe the connection between components and external clients, upstream and downstream data flows, where servers are located and how they are connected to the network.
Interviews and workshops	Enable more detailed discussions and identification of obstacles to migration. Once you've managed to collect basic information on applications, services and infrastructure, it is essential to get owners together to discuss the components in their area and how these can be best migrated.
Physical inspection	A physical audit of the boxes and cabling to identify serial numbers, available ports and wires connected to machines.
Technical documentation	System documentation, user guides, legacy plans and process documentation may be available for key systems. If not, you may need to create new documentation to fill in gaps.
Database management systems	Contain data on existing infrastructure and applications. Databases can be queried and used to generate reports. Some suppliers offer specialised tools and consultant services that are intended specifically for data migration projects.

Step 7. Validate the information collected

Validating the data collected is essential for ensuring its accuracy. Data collected from different sources may provide conflicting information. Some form of data cleansing, validation and sign-off is appropriate before this information can be used for planning. For example, data could be validated using a workshop approach, as described in Table 3.

Information can quickly become outdated. One of the problems when application discovery is spread over several months is that the information is already partly out of date by the time it is collected and presented to the business.

To control this, it is good practice to ensure that mechanisms are in place to identify and update information that has changed, without affecting the integrity of the data that has already been validated. It is also typical for businesses to impose a change freeze around the migration to minimise risk.

Tackling issues

Things rarely go smoothly during complex migration projects. It is always best to adopt a proactive and positive approach to issues that arise. The following issues, many of which were encountered during our migration project, are likely to be relevant to other such projects.

Finding suitable subject matter experts

Often, due to redundancies and staff turnover, key subject matter experts have been lost and documentation is not available for many systems. In this case, you may need to do more detailed research and investigation to find the information you need. If you can't find a nominated expert, try talking to alternative contacts and searching directories and intranets for additional information.

The value of such peripheral investigation should not be underestimated. I found it often enabled me to ask informed questions, which significantly improved the responses from available business contacts.

Dealing with cultural differences

Business and cultural sensitivity is essential when working across countries and cultures. Communication difficulties can be compounded when dealing with different language speaking areas — normal e-mail and phone channels may be less effective, so a trip to the region could be required.

It helps if you can find a contact in each region who can co-ordinate and escalate issues on your behalf.

Often data collection and migration projects are a preliminary step to eventually closing down or selling off business areas — with associated redundancies. Team members should be aware of these sensitivities when asking for information. During the data collection, try to find out if any plans have been made for selling off or outsourcing business areas and functions. This information will have an impact on how the migration is planned, and on what information you'll need to collect.

Third party suppliers and vendors

Components of a system may be managed or hosted by third parties — vendors, suppliers or partners connected to your business. You will need to identify relevant supplier managers to work with. Will moving your systems have an impact on the service they provide?

Information about licence agreements and maintenance contracts is important; licences may need to be repurchased and maintenance agreements closed or renegotiated, as part of the infrastructure and application migration.

Legacy systems may contain components with outdated licences. When the system is migrated, the components no longer function, as the licence has expired. At worst, new licences cannot be purchased because the vendor either is no longer in business or does not support the outdated version.

In some instances, systems are that are currently managed in-house may become candidates for outsourcing to third party suppliers who can host and manage the service more cost-effectively. This is particularly likely where a system needs to remain local or within the country of origin.

Disaster recovery considerations

Migration planning must take into account not only the impact of downtime on existing live systems, but also the backup and disaster recovery arrangements currently in place, and what needs to be put in place in the new infrastructure. This includes information on backup and recovery objectives, and data retention requirements.

Systems with disaster recovery need different planning and fail-over testing to move both instances.

You will need to work closely with business areas responsible for business continuity and disaster recovery planning, to ensure their information requirements are met.

Application criticality

Information on application criticality is essential for migration planning. It determines not only migration timing and resources required, but also fallback arrangements in the event that a system does not function after migration.

Since most application owners will tend to rate their applications as critical, it is important to arrive at a more objective assessment of application criticality.

You can use a number of key questions to determine how critical an application is. For example, how much revenue does the application generate? What transaction volumes does the application handle? How long can the business manage without it? What would be the affect on the business if the application were unavailable? What are the contractual agreements for service and availability for this application?

An application's criticality can also vary, depending on the time, day or season. For example, a finance application may be critical at month-end or year-end, but not at other times.

You can identify less critical applications by determining if the application is being phased out and is candidate for decommissioning, or if there are alternative applications that can provide the same service to the business.

Some migration specialists offer weighting systems, which provide a measure of criticality based on answers to these types of questions.

Identifying risks

During the project it is important to keep track of ongoing issues and risks. Project managers need this information to effectively manage the significant risks associatied with large-scale migrations.

For example, if there is not enough resource to complete the information gathering phase, you should identify this as a potential risk that could cause delays to the migration.

Various risk assessment methods and tools are available to help identify the most critical risks and activities.

Location of data centres

Physical distance between the data centres is important: there can be technical limitations (for example, when synchronising machines that are far apart) and timing issues (for example, physical transport of tapes).

Call centre staff and telephony equipment may need to be co-located, as there may be significant bandwidth and cost issues which make it unpractical to run such equipment from a remote location.

Systems in different data centres, even where running on the same hardware, may have different service level agreements. You will need to determine if these need to be harmonised.

Legal and compliance issues

Legal issues include current agreements with customers, as well as regulation in the country where the applications are currently hosted. A country's legislation may impose restrictions on movement of certain data or systems. Data protection laws may also differ.

In addition, there may be important tax implications for running systems out of different regions.

Data centre equipment

When consolidating many data centres, there are likely to be multiple copies of common items such as fax servers, routers and web servers. This needs to be identified early, since such elements may not be interoperable.

When equipment is consolidated, new licenses agreements may need to be put in place.

In addition to formal equipment contained in data centres or supported by IT, development teams and users may have unsupported equipment, usually sitting under their desks. As part of the data collection process, it is a good idea to maintain a list of all such equipment and their function.

Ongoing business projects

Migration in larger organisations is not a clear-cut process. There may currently be several projects under way, to deal with selling off, migrating or upgrading parts of the infrastructure. You'll need to identify what these systems are, who owns them, and when the upgrade, sale or migration is planned.

Identifying knowledge gaps

Migration planning needs to consider how systems will be supported post-migration. This includes the transfer of knowledge from old data centre staff to any new staff that must manage the systems. Good documentation is essential to support this process. You may need to identify where there are currently significant gaps in documentation.

Conclusion

Migration projects provide opportunities for technical communicators to be involved in the planning, data collection and communication processes supporting the migration. Large projects are likely to have shifting requirements and involve multiple parties, including business analysts, project managers and third party consultants — all of whom you'll need to work closely with.

As a technical communicator, you may only be involved in a small aspect of the data collection process described in this article. However, it is useful to have an understanding of how your efforts fit into the larger migration programme. Your role in facilitating good communication between business areas and achieving agreement can be important to the migration's success.

The data collection and planning effort will involve a significant cost to the business. A systematic approach can ensure that the impact on daily business activities is minimised, and help to facilitate the information-gathering process.