



An Insight company

Data center transformation: an application focus that breeds success

Introduction

Behind any significant data center transformation is often the act of migrating, relocating, upgrading, or consolidating applications and their accompanying ecosystems. Do it right and the business will thank you for its new opportunities to be fast and agile. Do it wrong and face the pain of unexpected downtime events, cost overruns and scheduling delays—not to mention the ongoing (and often vocal) dissatisfaction of the business and its users. Thankfully, you can increase your chance of success with your organization's data center transformation. This white paper outlines a detailed, application-centric protocol that ensures your applications (and the IT organization itself) experience smooth sailing prior to any significant change.

Data center transformation: the vision

Much has been written about the need for companies to transform their data center operations. And, with good reason. Data centers bound by legacy architectures and lengthy, business-as-usual resource provisioning have begun to feel the strain of an increasingly service-driven, public cloud world. The need to compete in the face of emerging shadow IT, big data, mobile, and social media requirements has wrought something of a revolution on the part of the typical corporate data center.

No longer is it enough to keep the lights on. Businesses are demanding—and getting—more from their IT organizations, whether that organization represents IT staff employees or external IT providers. Often prompting this data center transformation is the need for IT to become a true, agile service provider able to quickly and efficiently meet the changing needs of the business.

Many IT organizations have a general vision about how this type of transformation might occur. To get there, many may even have a good idea about the organizational, philosophical, and technological changes that need to be made. These might include:

- **A philosophical shift** from a reactive stance to one that more proactively drives business improvement.
- **A new organizational focus** on automation, responsiveness, and relevance that seeks to improve the business through fresh combinations of technologies, IT services, processes, and skills.
- **A holistic change** to IT that better aligns technology, services, and processes to the needs of the business and to the support of business-critical applications.
- **A technological shift** that will require core changes to the IT infrastructure. Many such changes may involve some form of data migrations, data center relocations, or data center consolidation. They may also involve the upgrade of hardware/software, re-architecting systems, and changing the way data is accessed and structured.

Such visions of a newly transformed data center inspire positive momentum, enthusiasm, and hope for a new world where the business and IT operate synergistically and in-step with one another.

But, moving from this vision and the above goals to their realization is often another story. How do you get from where you are to where you need to be? How can you be sure the transition will be as smooth and trouble-free as possible? How can you ensure your data center transformation projects won't be riddled with unplanned downtime, let alone schedule or budget overruns?

This whitepaper offers some key takeaways that will increase your chances for success.

Translating the vision into reality

While not a panacea to all the challenges that can await wholesale transformations, this white paper can help IT organizations successfully transform based on its answer to one key question:

How do you ensure the continuity and ongoing health of data center applications before, during, and after any significant transformation?

It should come as no surprise that significant planning is required before a data center migration, consolidation, or relocation project. What may be surprising, however, is how detailed the planning should be, especially as it pertains to the applications involved and their various dependencies.

The most effective transformations utilize a set of practices and methodologies that work time and again to ensure smooth transitions, minimal business disruption, and successful data center transformations.

The principles surround a detailed process called application dependency mapping (ADM). This process is performed at the initial discovery, analysis, and planning stages of most data center transformation projects.

What is application dependency mapping

ADM is all about understanding, at the application layer, what the dependencies are in order for that application to continue to function. In short, application dependency mapping lets you do four things:

1. Create relationships between interdependent applications, storage, network, servers, and databases.
2. Understand upstream and downstream relationships.
3. Create application maps that show the relationships of the applications to the infrastructure components that support these services.
4. Create an airtight, current state assessment of the infrastructure at a particular point in time.

This includes details about:

- ✓ The number of total applications supported in the current environment
- ✓ The number of servers and storage in the environment
- ✓ The number of networking components and their roles
- ✓ How the applications communicate internally and externally

Why ADM is important?

ADM requires you to ask key questions about applications before embarking on any significant data center transformation. It requires you to ask what other elements—whether hardware or software—does an application require in order to function properly? Maybe an application relies on one or more applications or database systems. It may rely on specific storage, servers, or network connections. Conversely, what other hardware or software elements, upstream or downstream from a given application, depend on that application to function?

ADM requires you to ask yourself, specifically, if you make a change, will it “break” anything else in your environment? Will it cause an outage?

IT organizations that practice configuration management or that do their best to maintain internal documentation about their operations and applications might think this information has already been adequately captured.

However, the process of application dependency mapping routinely uncovers previously unknown dependencies that could have caused a significant outage or issue during migration.

In some cases, prior dependencies might include forgotten legacy applications with hard-coded IP addresses or an application’s unknown latency requirements.

By committing to the process of application dependency mapping, an organization can expect to:

- Minimize the risk associated with a data center transformation project through more in-depth knowledge of the environment
- Increase resiliency and flexibility
- Smooth the transition to a new environment
- Ensure there are no unplanned outages during the transition
- Reduce time spent on the actual execution of a migration, relocation, or consolidation

The three-part process of application dependency mapping

For migration projects, it is best to follow a very specific methodology regarding application dependency mapping. This includes detailed deliverables and milestones to ensure successful outcomes.

While the details of this methodology are beyond the scope of this paper, it includes some of the high-level points for any organization seeking to move forward in their own data center transformation. This includes a three-part process:

- ✓ Part 1: Discovery of application dependencies
- ✓ Part 2: Analysis of previously discovered application dependencies
- ✓ Part 3: Planning a move, migration, or consolidation based on prior analysis

These parts can comprise up to 60-70% of a given data center consolidation or migration project. It becomes an example of the “pay now or pay later” philosophy. Time spent upfront will pay off when it comes to the actual execution of a move, migration, or consolidation. Done correctly, the process of application dependency mapping performed during the discovery, analysis, and planning phases can significantly reduce the time spent on an organization’s actual move or migration.

Part 1: Discovery of application dependencies

The first part involves a fair bit of detective work to uncover applications and their dependencies. This includes both manual effort and the use of automated discovery tools. The goal is to identify:

- All applications
- All servers, storage, and networks involved with those applications
- Devices that communicate with one another
- TCP ports these devices use for communications
- Processes that run on these devices
- Unique business requirements for each application, including recovery time objectives (RTOs), recovery point objectives (RPOs), service level agreements (SLAs), operating level agreements (OLAs), latency factors, security factors, outage tolerances, business criticality
- Details about how the applications work and how they are maintained and supported

Manual effort involved in ADM-based discovery

The manual sleuthing involved in identifying application details occurs with two key tasks:

1. **Review documentation:** Documentation that can help identify applications and their potential dependencies goes farther than operations manuals. In general, try to collect information from the following sources:
 - o The existing configuration management database (CMDB), if available
 - o Existing SLAs and OLAs
 - o Application design documentation
 - o Support tickets (if no documentation exists)
 - o Any documented custom or unique requirements identified for an application
 - o Any other documentation produced by the application team
2. **Interview individuals:** It's important to conduct interviews with a variety of individuals to supplement documentation. Figure 1 shows individual roles of value during this part of discovery.
 - o **Interviewing subject matter experts (SMEs) and application owners:** Interview individuals like SMEs or application owners to gain additional insights regarding such areas as: Outage tolerances or other unique and/or custom areas of knowledge regarding systems or system flows not otherwise documented.
 - o **Interviewing others of interest:** You may also be able to gain good detail about application-specific system operations, system maintenance, and support by interviewing system administrators, architects, managers, and end users.

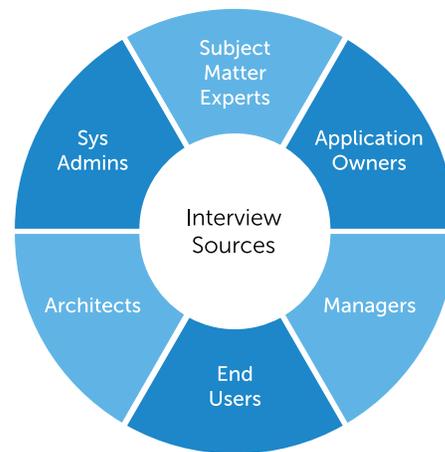


Figure 1: Internal interview sources can help research application dependencies

Automated tools used in ADM-based discovery

Deploying automated discovery tools in the infrastructure often augment, corroborate, or validate the manual information collected from documentation and interviews about application dependencies. Sample automated tools often involved during discovery include:

- Application Dependency Discovery Mapping (ADDM)
- Capacity planners
- Network traffic monitoring

Part 2: Analysis of previously discovered application dependencies

This second part of the ADM process draws conclusions based on the application dependency maps developed after the first part of discovery has been completed.

The purpose of this second part is to “group” applications and infrastructure together based on dependencies. See Figure 2 for a visual map with a sample grouping of applications and infrastructure.

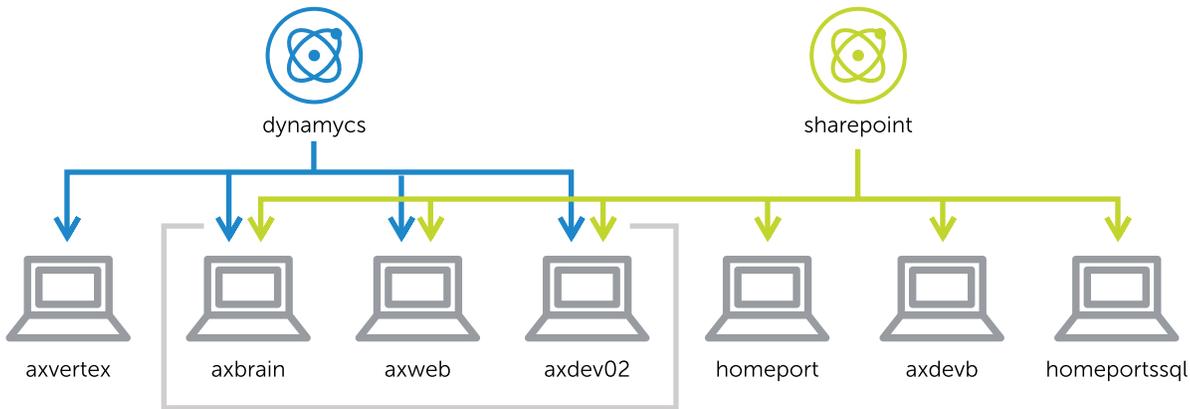


Figure 2: Grouping applications and infrastructure based on dependencies

The end goals of this grouping process are to:

- Create physical and logical application and infrastructure groups based on dependencies.
- Identify and plan for potential migration or relocation issues for each group prior to the event.
- Determine different types of move methods appropriate for each group, based on the group’s complexity, its criticality to business operations, or its specific hardware needs.
- Develop migration bundles based on a predefined grouping. Such migration bundles can be prescheduled and coordinated within Part 3, the planning stage.

Part 3: Planning a move, migration, or consolidation based on prior analysis

The third part, prior to actual execution of a migration, consolidation, or relocation is the process of developing in-depth checklists, schedules, and a detailed "runbook" of what will be moved at what time and in what order.

This planning phase involves much more than the components of an application dependency mapping exercise. However, the prior ADM grouping you identified in Part 2 will be used here to identify key "migration bundles" that must now be scheduled for migration in groups and stages in order to minimize risk to the business. Considerations here will include allowable outage windows per group, any testing required prior to or after the move, and mitigation of any potential issues before the actual event.

One example of migration bundles might be Application Group 1, which might have an unusually short outage window. Based on this outage window, it might require a different migration or move method than other groups. For example, if you had just one hour to move application groups and Group 1 couldn't complete the migration in that timeframe, you might have to create a migration bundle that moves Group 2 and Group 3 together during that window instead.

Figure 3 shows elements from a sample runbook developed as part of Part 3's planning phase.

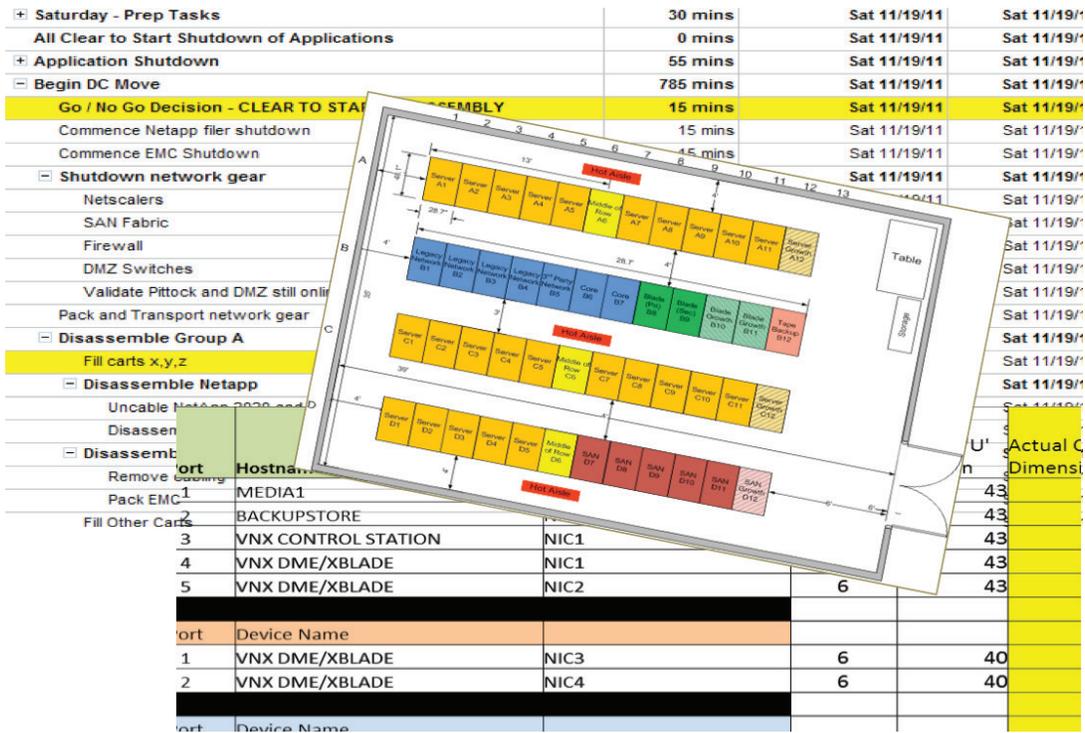


Figure 3: In-depth planning for a move or migration using groups and bundles



Follow Datalink online to receive the latest insight on data center technologies and practice.

twitter.com/datalinkcorp

blog.datalink.com

facebook.com/datalinkcorp

Application dependency mapping: turning vision into reality

As you can see, the prior three-part process for application dependency mapping is not for the faint of heart. It requires significant time spent discovering and analyzing key application details at the outset of any large transformation project.

It also requires in-depth program management skills to oversee the process and receive key signoffs at each step. While many organizations see the value in the in-depth process of application dependency mapping, they also may recognize that they simply don't have the bandwidth or IT staff to devote to such an exercise.

Beyond the prospect of hiring more staff, many turn to a solution provider to help guide them in this process. Many even rely on a service provider to be a short-term extension of their own IT teams.

Like you, a good provider wants your data center transformation to be as successful, cost-effective, and trouble-free as possible. Their goal is to help your IT team—and the business it supports—reap the rewards quickly.

If you'd like more details on our Turnkey Data Center Transformation Services, call us today and we'll discuss how we can assist your IT team.

About Datalink

Datalink is a complete IT services provider that helps companies transform their technology, operations, and service delivery to meet business challenges. Combining extensive experience, a full lifecycle of services, and a comprehensive approach to producing IT innovations that empower positive business outcomes, Datalink delivers success across cloud IT transformation, next generation technology, and security. For more information, call 800.448.6314 or visit datalink.com.



An Insight company

datalink.com | insight.com

© 2017 Datalink, an Insight company. All rights reserved. No portions of this document may be reproduced without prior written consent of Datalink Corporation. Datalink and the Datalink logo are trademarks or registered trademarks of Datalink Corporation.

DCT-WP-2.0.3.17