

**ON DISCOVERY:
FOR HYBRID IT STRATEGY
AND
DATA CENTER
WORKLOAD MIGRATION**

**GLOBAL
TECHNOLOGY
SOLUTIONS
GROUP**

**INFRASTRUCTURE
TRANSFORMATION**



WHETHER DEVELOPING A HYBRID WORKLOAD PLACEMENT STRATEGY, OR PLANNING A MIGRATION, **DISCOVERY** IS FUNDAMENTAL TO SUCCESS.

Inventory accuracy for the purposes of server to application mapping and affinity group development is far more elusive than many firms expect.

In a research paper released June 30, 2017, Gartner analyst Dave Cappuccio focused the attention of his considerable readership on one of the most important challenges facing I&O leaders as they look to take advantage of what the cloud has to offer. In sum: cloud has come into the mainstream, but it is not what 451 Research calls the “best execution venue” for all workload. Determining which applications and services are most appropriate for the cloud, and then planning to move them, is not an academic undertaking. Rather, it’s all about the details.

This paper offers some of our thoughts on discovery and information gathering. These ideas apply equally to- in fact, they have their origin in- the significant planning required for a data center migration. Inventory accuracy for the purposes of server to application mapping and affinity group development is far more elusive than many firms expect. At the end of the day – it’s the old “you don’t know what you don’t know” that can derail what might otherwise have been a successful move wave implementation. Our approach is designed to minimize the unknowns and associated risks.





GAP IDENTIFICATION AND REMEDIATION

The good news: many firms have tooling in place which is operational, albeit in a manual mode and with limited topology discovery.

In many cases, we believe that current tooling is capable of providing the data required to (1) identify all assets, and, (2) perform some degree of interdependency mapping. Granted, we need to remediate some gaps in the current deployment to overcome what appear to be “data breadth” issues – but many firms have tools which are installed, operational, and producing outputs.

We feel these issues can be resolved far more easily than starting all over again with yet another tool.

In situations such as this, our recommended approach would be

Identification of Gaps: we see various symptoms, including



- accurate but incomplete data;
- accurate but high-level data;
- inconsistent data across an asset base;
- stale or potentially non-existent data across an asset base

Most commonly we see two key issues with the data:

(1) Breadth of the data – only part of the asset base is being reported, and,

(2) Depth of the data – only part of the specific data on a given asset is being reported.

If we identify and remediate these two most common gaps, the CMDB data can then be used effectively in the application dependency mapping portion of our methodology.



ADDRESSING BOTH DEPTH AND BREADTH OF DATA

Fixing “Breadth of Data” issues: only part of the asset topology is being reported upon. **Remediation actions include** leveraging other asset identification methods, and sources at the IP level.

Methods include:

- DNS records (dynamic and static) / Active Directory
- Mainframe IP SMF records (source and target IP are of interest to us)
- NetFlow data
- Network “Sniffer” data
- Firewall rules
- Router-based static NAT rules
- Use of “Near Free” tools like SolarWinds

We’ve successfully employed all the following discovery methods: Passive, Agentless, Agent-based, and our choices at a given client will be governed by what makes sense.

Fixing “Depth of Data” issues: likely to be the second gap encountered after a client finds they are reporting on the entire asset base.

Remediation actions include:

- For purposes of the data center relocation program, confirm that the depth of data issue deals with information gaps that are required for the data center program
- Evaluation and appropriate updating of configuration within the legacy system
- Identify isolated test or development server; experiment with configuration and agent if utilized on this server

Identify use case were depth IS being reported accurately



THE RIGHT MIX OF BOTH TOOLS AND KNOWLEDGE

As we have stated many times over the years, there is no such thing as a “silver bullet” tool. No tool is capable of providing what years and years of institutional knowledge can provide. Our approach and methods enable us to leverage the “jump start” capability of a tool and then validate and enhance the output with focused interview/workshop approaches to identifying the true – and sometimes hidden – dependencies.

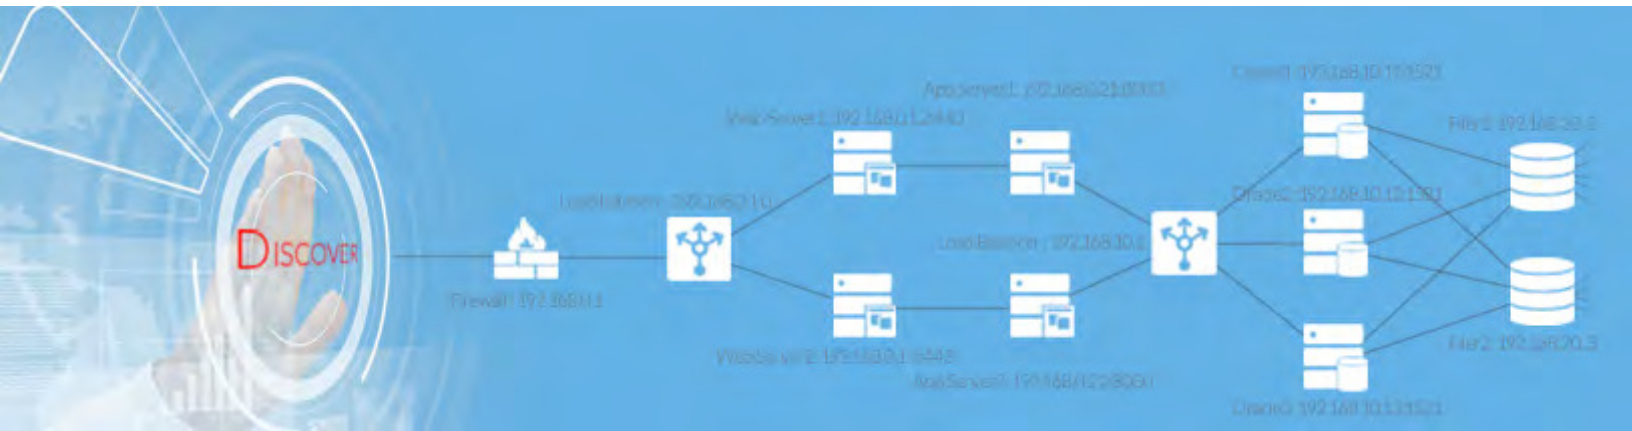
The risk of a “tool only” approach is that you forgo the priceless institutional knowledge to provide checks and balances. You also run the risk of being presented with an affinity group (the smallest unit of work that must be moved at the same time) that looks like it resembles 50% of your application portfolio. A majority of application move groups’ affinities looking the same is a red flag.

Our methodology has been proven in situations where discovery and interdependency mapping was produced manually, and in other data center events where a full set of tool-based interdependency mapping tools were refined with our workshop-based validation approach.

When considering augmenting your tool set, we advise against the use of specialty tools with which the team is unfamiliar, and advise against permitting independent change processes, or a process that is too narrow in scope to capture all changes that may have an impact.

Mainframe Considerations: Most CMDB tools will do little for the mainframe. We’re not saying they are ineffective on mainframe, but CMDB tools are definitely designed for distributed environments. GTSG would recommend a review of TCPIP SMF records to identify distributed systems that touch the mainframe and a review of application-related SMF records to assess data dependencies and such. Scheduling tools can also be utilized to research and ascertain system-to-system dependencies.





GTSG VALUE ADD: APP DECOMP

GTSG provides an additional level of **Interdependency Analysis** to be used at minimum for your critical applications. We call it Application Decomposition or “App Decomp”.

We employ a tested and proven technique of decomposing the application suite utilizing a refined series of detailed questions and workshops to identify hidden dependencies.

As an example, no tool will determine if a particular application’s DB access is “chatty” in nature. This characteristic isn’t apparent when a data source is close to the application end user (workstation-based ODBC applications as an example), but we have seen it surface after a data source moves 1,000 miles away as part of a Move Event.

Our App Decomp Analysis will expose such a characteristic during Discovery and avoid the adverse business impact of learning it after a move.

Additional Level Setting Topics: These areas are pertinent to Discovery and Information Gathering and we cover these early in an engagement:

- Confidence in the data you’ve based your business case upon.
- Identification of the trusted source(s) of data.
- Current state of application performance baselines and latency concerns.
- Level of understanding of your software licenses and service provider costs, and the contracts that govern them.

SUMMARY

Lessons from some of the decade's largest and most complex cross-country migration events:

What worked was:

- Development and refinement of a repository that was kept up to date and accurate
- “Trust but verify” a small set of tools, and not assume a tool is the silver bullet
- Workshops with application owners and infrastructure support teams to validate a base set of information collected by the program team
- Make the application teams accountable for accuracy of Point of Departure diagrams and details

What didn't work:

- Relying exclusively on existing CMDB or accounting information, or
- Expecting that tools would be able to automatically update the repository

What to watch for:

- Don't fall into the trap of spending too much time and money analyzing tools
- Keep the amount of data collected to the bare minimum to avoid “techie nirvana”
- Interfaces, protocols, data flows and chattiness characteristics are key to know

A Step Further – Our experience is that exposing end of life and out dated applications and infrastructure requiring upgrades during Discovery provides an opportunity to reduce inherent risk prior to the migration and eliminate lengthy troubleshooting and back-outs

What worked:

- Establishing a rigorous process to review and remediate all applications prior to migration. Items included were hard coded-IP's, MAC addresses, FQDNs, device and user names, passwords embedded in connection streams
- Using a migration factory team armed with tools to scan code for discrepancies
- Uplifting or migrating applications to new infrastructure including moving to a virtualized environment or at a minimum replace aged, EOL hardware

What didn't work:

- Not tracking responsible teams and providing a scorecard that drives accountability
- Asking applications teams to identify issues without the assistance of a factory type approach. The application team's priority is driving business by adding functionality

What to watch for:

- Ensure all requirements are identified prior to scanning code
- Get to a highly virtualized environment prior to migration – this will be one of the biggest cost savers and will substantially increase availability



Thank you for your interest in our observations.

For 30 years, GTSG has worked exclusively on Data Center transformation and optimization.

Please reach out if we can help you in any way.
Thank you.

**GLOBAL TECHNOLOGY
SOLUTIONS GROUP**

T 877 467 9885

F 877 225 4084

W gtsg.com

E itservices@gtsg.com

Data Center Strategy and Migration

Strategic Approach

- Business case development
- Transition planning
- Technical modeling
- Non-disruptive execution

Application Analysis Methodology and Tools

- Decomposition
- Affinities
- Wave planning

Project Leadership

Implementation Subject Matter Expertise

Infrastructure Support Services

Managed Services

- Multi-platform including DB & MW
- Service-level based or FTE-based
- Architecture, administration, programming, systems management
- Remote or Onsite

Project Based Services

- Platform upgrades
- Workload migrations
- Implementation services
- Consulting and Assessment (performance, DR, HA....)
- Project Management

Infrastructure Transformation

Transition Services

- Insourcing/Outsourcing
- Knowledge transfer and interim support
- Application migration
- Service management design

Disaster Recovery Design and Implementation

High Availability Design and Implementation

Application Assessment and Deployment

- Reference Architecture
- Infrastructure Alternatives/Recommendations
- Implementation/Migration

Infrastructure Optimization

Architecture Assessment and Design

Server Virtualization/Consolidation

Storage Optimization

Data life-cycle management

- Tiering
- Standardization/Automation

Application Decomposition Application

Re-design/Remediation Performance

Management and Tuning Latency

Analysis and Consulting