

Improving Change Impact Analysis in Data Centres

David Cuthbertson, Square Mile Systems

Predicting the impact of changes in data centres is becoming more difficult with the increasing complexity of hardware and applications. This white paper covers how the technique of combining data centre and service management knowledge within AssetGen Planner can drastically reduce the time, risk and workload on teams or managers, when analysing the impact of planned or emergency changes.

What do we mean by change impact analysis?

For this white paper, the term impact analysis means identifying the potential impact of disruption caused by a change to one or more components in the data centre on other components, including software and services.

To identify change impact, many organisations typically rely on a combination of local knowledge and data held in disparate systems of varying accuracy. Just to make sure, there is also manual tracing of connectivity and reverse engineering by phone calls, workshops or email communications. The end result is that projects and operations teams have to devote significant time and resources to understand existing infrastructure.

Figure 1 shows that impact analysis can take place at many levels with most data centre activities initially looking at hardware impact, but needing to also understand service impacts for effective change control.

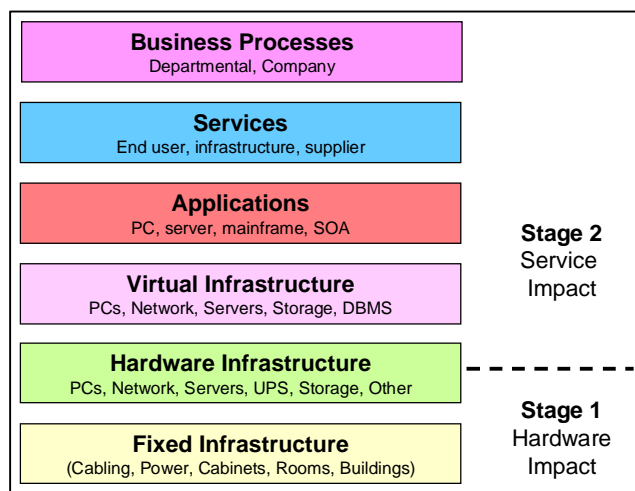


Figure 1 - Impact at different levels

Stage 1 Identify Hardware Impact

The basic process steps for change impact analysis are shown in Figure 2. We have the option to follow this approach for a full impact assessment, or to bypass some of the steps to suit other needs.

1. Start Point – We choose one or more starting points which could be anything physical from a site, to a cabinet, down to a specific port on a switch or PDU.
2. With the start point as a filter, identify a range of selected device types such as SAN boxes, power strips, servers, etc.
3. Trace the connectivity from those devices, looking for target devices such as critical servers.
4. If necessary trace again for onward connections, e.g. initially trace all edge switches connected to a core switch, then trace everything connected to

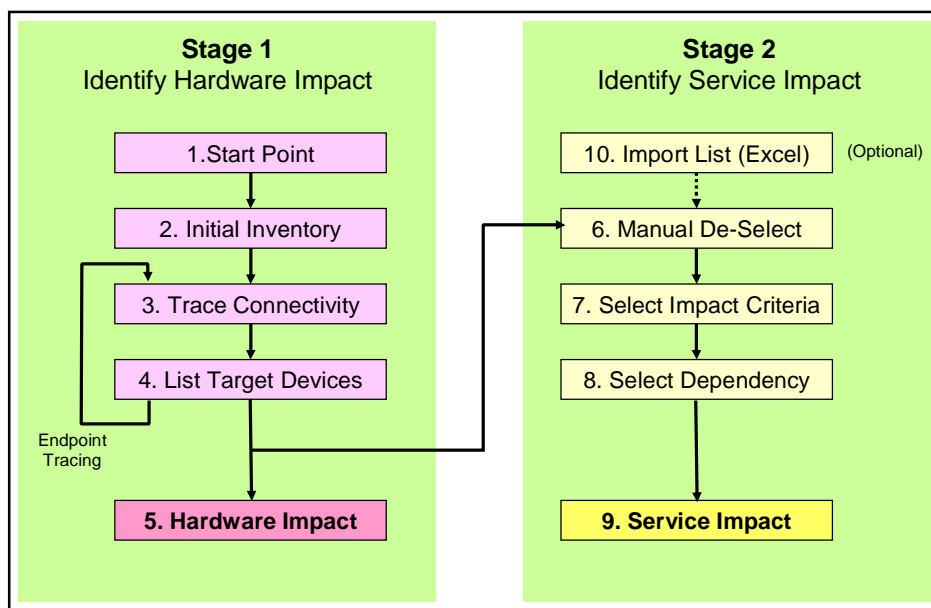


Figure 2 - Change Impact Process Flow

- those edge switches, looping a number of times to cover more “hops”.
5. Produce a list of hardware (including virtual servers) which we are interested in. Export this list if necessary, so we can insert it into a spreadsheet, or change form.

Within a few seconds we have identified everything of importance connected within our infrastructure and summarised it into a list. Case studies later in the white paper cover typical data centre tasks or situations.

Stage 2 Identify Service Impact

The next step is to determine the impact on software, services or business processes. We typically make use of service mapping techniques (Figure 3) as recommended within the ITIL and ISO20000 frameworks, or you can also use other dependency types such as batch processes, data flow, system architectures, etc.

6. From the hardware list, de-select devices which are known to be unimportant or irrelevant to the impact analysis.
7. Choose the level or type of impact - it could be services, software or business processes - you want to assess against.
8. Within AssetGen we provide the default ability to assess against service relationships, but you can also select other dependency types such as batch process, governance, departmental, etc.
9. Produce a list of services impacted for inclusion in project and change documentation.
10. Optional – Take an external list of hardware, software or other configuration items (CIs) and present them as an input to the service impact stage. This is ideal if you are upgrading multiple platforms as part of a technology refresh, or have a list of un-supported software and want to know which services are impacted quickly.

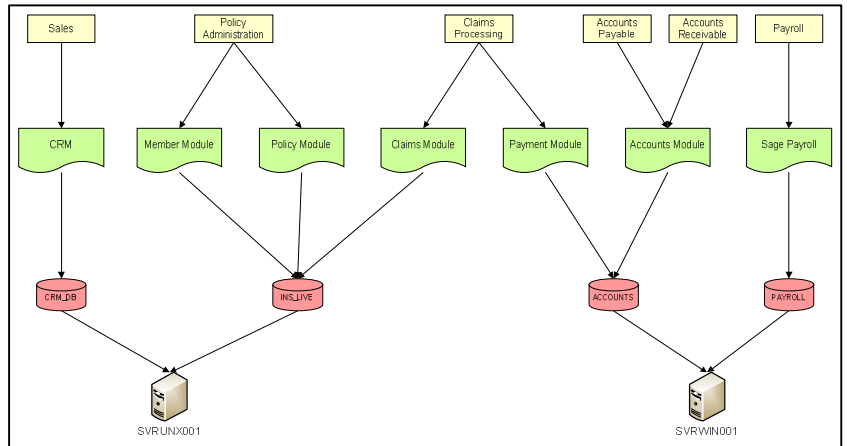


Figure 3 - Mapping Components to Services


In Stage 2 we have identified the service impact(s) across multiple devices, leveraging data held by other groups such as the service management team. We can bring instant service awareness to data centre and technical project teams without them needing to log into the service desk CMDB, or spend hours trying to find the person who “knows”. If this takes a few seconds to complete, how does this compare to the current method of identifying service impact?

Example 1 – Identifying Services Impacted by OS Patching of Servers

This is simple and quick in AssetGen Planner as we probably already have a list of servers that are targeted for patching, the difficulty is identifying whether one or multiple services are at risk with each patch.

- a. Go straight to a Stage 2 - Service Impact
- b. Import the list of servers from a spreadsheet or text file.
- c. Select services as the target impact and perform a multiple service impact analysis.
- d. Print or copy the results showing each server and the services impacted (see screenshot).

Follow the same steps for ten, a hundred or a thousand servers – it doesn’t matter if they are virtual or physical.



Example 2 - Services impacted by a PDU power down in the data centre

It doesn't take long to switch off the power to a PDU, but it does take a long time to identify everything that could be impacted at a hardware level, as well as at a service level. The key output is a list of affected servers at a hardware level, from which we can then determine service impact. Why is this difficult in practice?

We have to cope with not just servers connected directly to the PDU or through power strips, but we also have to consider infrastructure devices such as LAN and SAN switches which may also be connected at a power level. These switches then need to have their connectivity traced to any servers (which may not be directly affected by the PDU power). Not a simple task, but a common one when undertaking maintenance activities, or installing new electrics or cabinets. What about virtual servers or logical partitions sitting on the hosts? All covered by AssetGen Planner as part of the analysis function to save manual searching and interpretation.

The steps in AssetGen Planner would be;

- Select the PDU as a start point to begin Stage 1 – **Hardware Impact**.
- Trace all power strips connected and directly connected hosts (mini, mainframe, etc.)
- Trace all power connectivity to hosts and LAN/SAN switches connected to affected power strips.
- Trace all data connectivity from affected LAN and SAN switches to hosts.
- End Stage 1 with a list of hosts directly or indirectly impacted by power failure.
- Start Stage 2 - **Service Impact** - Using the host list, de-select any hosts that should not be processed (development, UAT, etc).
- Select business impact as the impact criteria.
- Finish with list of affected services and the hosts involved.

Undertaking this type of impact analysis is not simple, especially if you don't have the data in a consistent form to begin with. Even if you did have the data in spreadsheets, it would still be a major task to trace and summarise the data to produce the end result. When a generator fails, a UPS becomes unbalanced, or an electrical fire is suspected, how long will it take to determine the impact of a PDU power down?

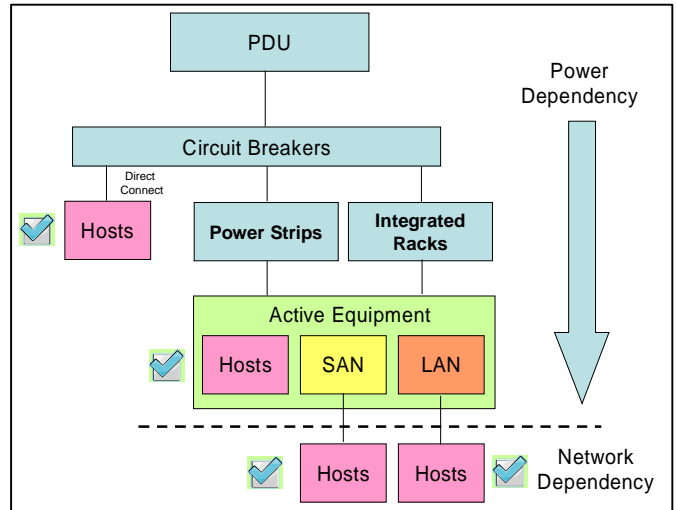


Figure 4 - PDU Power Down

Example 3 – Implementing a Change Freeze to Support a Critical Service

To ensure that a critical service is available at a crucial time, one management technique is to apply a change freeze to all components. AssetGen Planner can help quickly identify the supporting infrastructure components for a range of servers, so that no one forgets they are not to be touched without prior approval! – extending change controls beyond the service desk CMDB down to the infrastructure.

- Select the critical servers as the start points.
- Trace data connectivity to LAN, SAN and backup devices.
- Trace power connectivity to power strips and circuit breakers.
- Print off list of supporting components.

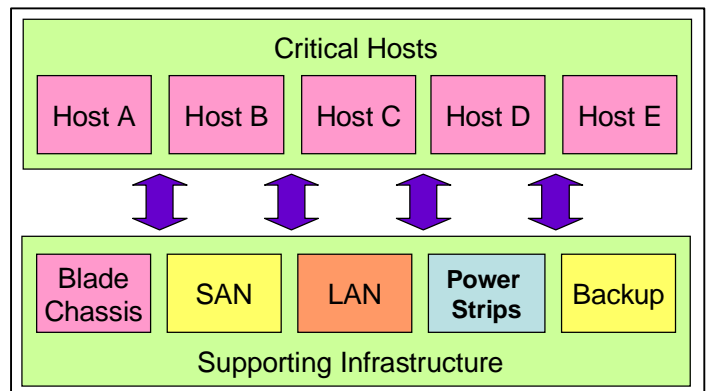


Figure 5 Where to Apply a Change Freeze

A similar approach could be taken for underpinning SOX controls on the financial reporting systems.

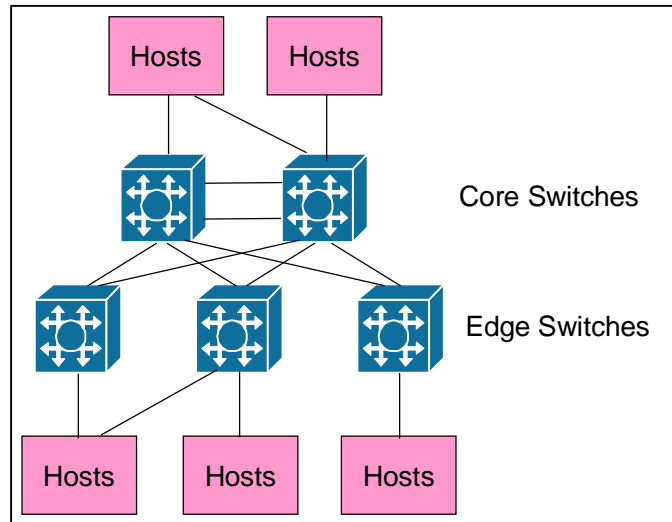
Example 4 – Rebooting Existing Core LAN Switches

Hardware or software upgrades, incident resolution or maintenance activities may all result in needing to reboot one or more core switches. Although resilience is assumed, in practice the risk must be assessed as not all devices have multiple network links.

The diagram on the right shows that some hosts may be connected to the core, while others are connected through edge switches.

The steps in AssetGen Planner would be;

- a. Stage 1 - Select an inventory of the room to find all core and edge switches.
- b. Trace all switch connectivity looking only for host devices.
- c. The result is a list of hosts and their status.
- d. To understand the risk to production services, then run a service impact analysis as a Stage 2 task.



The speed at which you can identify, select and trace 1000s of connections makes AssetGen Planner a must for project and change control, where networks are outsourced and the supplier is not responsible for knowledge beyond the switch port. If you think this same task could be achieved using NMS systems and auto-discovery then think again, you'll have lots of manual work ahead.

Making It Happen

If you have found this white paper interesting and practical in nature, then you will also appreciate that AssetGen Planner relies on supporting information from the AssetGen Connect and AssetGen SysMap applications. Key to the AssetGen approach is that information from other systems can be leveraged to avoid lots of manual collation, interpretation and summation. The problem is not a lack of data, just too much in different formats and spread across disparate systems.

Start small, focus on the critical systems, embed new working practices with each major technology refresh and remember that continuous service improvement starts as soon as you challenge the existing methods (or lack of) to see if there is a better way. Where you have a green field site, then it would be a missed opportunity not to have AssetGen Planner help you with the transition planning, the detailed change control and the subsequent in-life support. We didn't even cover the other 3 parts of AssetGen Planner – automated Visio diagrams, searching for cabinet space and capacity reports!

The Author

David Cuthbertson is a founding director of Square Mile Systems, a UK computer services company focused on configuration management solutions. He is also a director of AssetGen Limited which develops software applications for infrastructure management. He is a regular industry speaker at BCS and ITSMF events on best practices and the application of configuration management techniques to ICT and data centre infrastructure. Previously, he was the chairman and a founding member of the BCS Service Management Specialist Group (SMSG).

Square Mile Systems
3 Church Street, Cirencester
Gloucestershire, United Kingdom GL7 1LE
Tel +44 (0)870 950 4651 Fax +44 (0)870 751 9268
www.squaremilesystems.com

