GÉANT
Networks • Services • People

Service Prioritisation as part of a DC Continuity Plan

Best Practice Document

Produced by the CSC/FUNET-led AccessFunet working group

Authors: Janne Oksanen (CSC/FUNET), Janne Niemi (CSC/FUNET), Hannu-Pekka Poikonen (Univ. Of Aalto), Kaisa Haapala (CSC/FUNET), Tuukka Vainio (Univ. Of Turku

March 2016

Document No: GN4-NA3-T2-FN1.2
Version / date: v1.0/March 2015
Original language : Finnish
Original title: Palveluiden priorisointi osana konesalin jatkuvuussuunnitelmaa
Original version / date: v1.0/March 2015
Contact: accessfunet@postit.csc.fi

The work has been carried out by a CSC/Funet led working group AccessFunet as part of a joint-venture project within the HE sector in Finland.

The research leading to these results has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 691567 (GN4-1).

Table of Contents

1 Introduction

Campus networks have become very important and even critical for the functioning of the entire campus. In addition to basic network services, such as IP addresses, DNS services, and dynamic address management (DHCP), the network services offered include e-mail service, printing services, storage services, and various information systems needed by the study administration or financial administration, such as study attainment registering systems, working time tracking systems, or invoice payment systems. The services are provided from either physical or virtual servers, while the telecommunications are handled with, for instance, routers, switches, and firewalls. The data centre lives at the heart of the campus network, housing all these services critical to the campus network and the hardware they require. A data centre typically houses several different pieces of equipment.

In order to operate according to its purpose, a data centre requires some kind of cooling system in order to prevent the temperature in the data centre from rising too high for the hardware. In addition to a cooling system, a data centre requires a powerful electricity supply system able to supply the correct amount of wattage to all hardware. These two are the most critical systems in the data centre. Should a disruption occur in one of them, you must know in advance what actions to take. A business continuity plan must be prepared in advance for these kinds of major incidents, detailing exactly what to do.

This Best Practices document focuses on the prioritisation of services and service platforms as part of the business continuity plan, when an unforeseen major incident occurs at the data centre, such as a power outage with no reserve power immediately available, or a malfunction in the cooling system that cannot be repaired right away. During an incident such as these, a great many services may have to be shut down. This document does not list all issues in full detail; it merely provides examples that provide a good starting point for thinking about the documentation, instructions and operating processes in your own environment that will be part of your business continuity plan.

2 Documentation

Documentation of the data centre's entire technical environment and its processes is the alpha and omega of everything. We recommend using a some electronic system as the documenting tool in order to make updates as easy as possible. Electronic systems may also support linking a single entry to several different places, eliminating the need to update the same entry many times over. Issues can be documented in, for example, a Configuration Management Database (CMDB) system; these are available as both open source and commercial applications. Issues can also be documented into several different systems, but in that case, you must ensure that information is updated in all of the different locations affected.

Documentation is also useful during incidents, for instance in the case of a hardware malfunction. Deploying the backup service will be faster in a managed and documented environment than in an unmanaged environment. Furthermore, in a managed environment, the documentation includes the information on what needs to be done when a service requires additional resources, such as more hard disk space, processing capacity, data transfer speed, power, or cooling.

Use standardised terms and concepts in the documentation. This makes it easier to describe them for FitSM/ITIL-compliant processes [REF 1 and 2].

Document the operating environments and their service provision possibilities, including any limitations. The information system architecture often includes one or more fronted servers, with the data located on the background on one or more database servers. In this kind of an environment, it is important to know the interdependencies of the servers so that any errors caused by the network and the servers can be managed and the service brought back up again fast, so determine and document the interdependencies of all hardware. Drawing pictures may be of a great assistance in this.

Document the redundancies of the network and the services, which equipment or systems act as backups for each other, and where the SPOFs (Single Point of Failure) are. Also, document what manual actions need to be taken, if it was not sensible to automate everything. Without comprehensive documentation on what equipment, systems, and server rooms are available, it is almost impossible to successfully manage things, particularly during major incidents.

The documentation must also include information on who makes the decision on the measures to be taken, who will carry out the decided measures, and how issues are communicated to the end-users. Up-to-date contact information must be available to all personnel, as well as information on their deputies.

Divide the data centre environment into smaller parts/categories in order to more easily perceive what needs to be documented. See below for an example in which a data centre environment has been listed at three different levels: hardware, network and server/service.

2.1 Data centre room infrastructure

Document the available data centre rooms, their physical size, load-bearing capacity, electricity supplies (phases and powers), cooling machines, and reserve power supplies. Name the rooms, and draw up access instructions for incidents, including information on those individuals who are allowed access into the rooms. What is the room's maximum thermal load and its cooling power. If remote management connection of the services fails, draw up instructions on how to access the data centre

rooms and thereby access the management interface of the service. Appendix 1 lists more examples of things that need to be documented.

Disruptions in the data centre room infrastructure may cause an extremely wide-scale impact on the services produced in the room. The disruptions may be caused by various malfunctions, but planned maintenance also significantly raises the risk level of the service continuity. You can attempt to prepare for cooling system malfunction or outage in the external electricity distribution network through hardware prioritisation and a shutdown plan. However, a malfunction in the data centre's own electricity supply may cause a power outage directly affecting the IT hardware in some of the centre or perhaps even in the entire centre. Recovery from such a situation may require special measures, for instance making sure that the devices do not start automatically in the wrong order. The uncontrolled shutdown of hardware may cause problems to both the hardware and software.

2.2 Network infrastructure

Document all routers, switches, firewalls, and any other telecommunications equipment in the operating environment. What equipments are they, what services are produced with them, and how and with what are they managed and monitored? What is their capacity? Is any kind of automation used in any of the settings, etc.?

One outcome should be a list of all telecommunications equipment, their features, interdependencies, and management. Here, a list can also refer to a hardware register.

We recommend labelling the equipment with uniquely identifying codes so that they can be located quickly in the data centre when necessary. With regard to firewalls, you should also document their rules and the rule maintenance process. It is important to know the ports and protocols required by network services, particularly if a service is moved from one data centre to another due to an incident, for instance.

With regard to network connections, we recommend documenting the network topology and the physical connections between hardware. The network documentation should include the address spaces used for each service. It should also include descriptions of how the management connections, VLANs, and routings have been set up. See Appendix 2 for more examples of things that need to be documented.

2.3 Server and service infrastructure

Document both the virtual and physical servers and their hardware properties: CPU, amount of memory, number of power supply units, are there internal hard drives, etc. – corresponding information as for the network hardware. Document the locations of the servers: which data centre and rack are they located in. Document the services used and their dependencies on other services, for instance, the dependency of a service on storage servers. Also, include whether the service runs on a virtual server in the documentation. If yes, which server/what is the name of the server? If the service requires its own, physical server, document the dependencies on other servers and network disk shares that the service requires.

Document those services that are most critical for telecommunications: DNS service, DHCP, NTP (its physical servers, where are they located, what kind of redundancy is there for the hardware and connections)? See Appendices 3, 4, and 5 for more examples of things that need to be documented.

As part of the documentation process, you should try to draw up a list of the services in the order in which they can be shut down and restarted. Draw up separate instructions for each service on how they can be shut down and restarted, and how their operation can be tested.

Once you have finished the documentation, remember to communicate so that the required parties can find the documentation and instructions. Also, test that the persons are able to log into the necessary systems.

3 Prioritisation

Once you have documented the interdependencies of the hardware and services in the data centre, group them and draw up a prioritisation within each group. The groups could look like this, for instance:

1. virtual servers and standby servers
2. virtual platforms
3. physical servers
4. backup systems
5. disk systems and database servers
6. critical basic network services (DNS,DHCP, NTP, AD,...)
7. telecommunications equipment

Next, insert the hardware into the groups on the list. Remember to check the dependencies from the documentation so that, for example, fronted servers are shut down before the backend servers. Prioritise the hardware within each group. At this stage, consider which device is the most important in each group, and assign it a higher score than for less important hardware. If several data centres are in use, you should also take into consideration situations where a service can be moved to a backup data centre.

The following sample list contains virtual servers "testserver" and "customerservice2" that run on the Blade1 virtual platform, and "customerservice1" runs on the Blade2 platform. These servers must therefore be shut down before their platforms. In the example below, the order has been prioritised so that the smallest adverse impact is caused by shutting down "testserver" first and "customerservice2" second.

1. virtual servers
1.1. testserver
1.2. customerservice2
1.3. customerservice1

2. virtual platforms
2.1. Blade1
2.2. Blade2
3. physical servers
4. ...

You should consider the prioritisation from different perspectives, but you should only choose one or two "grand designs" and draw up individual prioritisation lists for these.

Prioritisation criteria may include electricity consumption, SLA with the customer, number of end users, the impact of the service on other, more critical services, etc. It is easy to come up with different perspectives, and you should apply them, particularly when prioritising within the groups. See below for two examples where "Play for time" and "Shut down the entire data centre" have been chosen as the grand designs.

3.1.1 Example 1: Play for time

In this case, a major incident has occurred at the data centre, for instance, a power outage the duration of which cannot be precisely known, although some sort of information may be available on its duration or guesses can be made, for instance based on the energy company's bulletin. If reserve power is possible at the data centre, but there will be a certain delay until the agreed reserve power supply is delivered on site, it is not sensible to shut down the entire data centre. In this situation, you should prioritise the services that consume the most power and/or service level agreements with the customer that allow shutting down these kinds of services among the first batch.

Once some of the services have been shut down, monitor the status of the UPS and cooling and assess whether you need to continue shutting down services or can you wait for the power to come back on or the reserve power supply to be connected.

3.1.2 Example 2: Shut down the entire data centre

In this case, the incident at the data centre is so major and severe that all hardware must be shut down. Keep a close track of the previously prepared prioritisation list and shut down the services accordingly. It may be possible to shut down certain services simultaneously, but in these situations it is very important to keep track of the dependencies in order not to skip too far on the list and shut down a platform still running services.

Prepare individual prioritisation lists for each data centre. On the list, you can mark the dependencies of the systems with colours (used in the table below (red and green ones)) or by adding a separate column for the purpose and making clear entries in it. Remember to update the documentation regularly. Include documenting as part of the service processes, so that it is naturally linked into other operations. For instance, when setting up a new service, you should check that the documentation is up to date and, for instance, access management works correctly.

See below for an example of a prioritisation table. In the example, you can follow the list systematically from one group to another (the numbers are in an ascending order), but in a real data centre environment you may have to jump between groups. It is thus important to consider how to mark the prioritisation on the list so that it is easy and quick to read and follow in various situations.

GROUP	PRIORITY	SERVER	DOWN	UP	ADMIN	MANAGEMENT	GUIDE	LOCATION	POWER	PHASE	FUSE	UPDATED
1.0 Virtual and standby servers												
virtual	1.0	virt1-srv			admin1	Wmware/ 192.168.1.10	service1-guide.txt	Blade1	100W	(3-phase)		6.2.2015
physical	1.1	srv1			admin2	Console/ 192.168.1.7	see. Server manual	Rack Y1.5	500W	Phase I	16A	4.8.2013
physical	1.2	srv2			admin2	Console/ 192.168.3.45		Rack Y1.6	300W	Phase III	16A	8.4.2014
virtual	1.3	virt4-srv			admin1	Wmware/ 192.168.5.7		Blade4	100W	(3-phase)		6.2.2015
...												
2.0 Virtual platforms												
Blade 1	2.0	blade1			admin1	Wmware/ 192.168.8.4	see blade guide	Rack X3.1		(3-phase)	32A	6.2.2015
Blade 4	2.1	blade4			admin1	Wmware/ 192.168.1.1	see blade guide	Rack X3.4		(3-phase)	32A	6.2.2015
Blade 2	2.2	blade2			admin3	Wmware/ 192.168.5.1	see blade guide	Rack X1.1		(3-phase)	32A	6.2.2015
...												
3.0 Physical servers												
server1	3.0	servu1			admin4	Console/ 192.168.10.5	servu1 guide	Rack Z1.8	300W	Phase I	10A	8.4.2014
server3	3.1	srv3			admin5	Console/ 192.168.10.6	srv3 guide	Rack Z2.1	250W	Phase I	16A	8.4.2014
server2	3.2	servu2			admin1	Console/ 192.168.10.3	see srv-guide	Rack X3.7	300W	Phase II	16A	8.4.2014
...												
...												

Figure 1: Example of a priorisation list

The table columns "DOWN" (shut down) and "UP" (started) can be used as assistance when going through the list. By ticking the box of each service, you can quickly see the overall situation, and monitor and prepare for the next necessary measures.

Once the incident at the data centre is over, starting up the services is quick when you can follow the prioritisation list used in the shutdown process backwards. This also means that the services are started up in a controlled manner. Once the services have been started up, you must remember to test that they operate correctly. The users should also be informed of the incident, at the latest when you know the cause of the incident – if you have been unable to inform them already during the incident.

Place the prioritisation lists and documentation in a place from which the data centre staff / administrators can find them quickly. If you use an electronic system, ensure that it will also work during a major incident (e.g. the entire data centre has no power). If the list is relatively short, you can print a paper copy at regular intervals. Also remember to train the data centre staff in the performance of the necessary measures.

4 In case of an incident

When/if a major incident occurs at the data centre, keep calm. Attempt to determine the cause of the incident and estimate its possible duration. Inform the data centre staff and the parties affected. Prepare to move services to a backup data centre if possible. If the incident continues, make a decision on which prioritisation list you will put into effect. Follow the prioritisation list and start shutting down services according to the plans.

Once the incident is over, start the services back up by following the prioritisation lists.

Be careful of exceptions / temporary solutions. They can easily remain alive in the network, where they can create new dependencies that might be difficult to eliminate.

APPENDIX 1: Room, racks, cooling, electricity, reserve power, fibres/copper cables, access, and contact information

See below for a list of things related to data centres that should be documented.

- Name of the room and its contact information
- Size of the room (length x width x height)
- is the floor raised?
  - if so, what is its load-bearing capacity
- Equipment racks
  - floor plan
  - unique identifier of the rack
  - rack size (width x depth x height, also height in RU)
  - load-bearing capacity and weight
  - where in the room is the rack installed
    - in the case of a new installation, follow the floor plan
    - the plan includes, for instance, the physical location/placement, electrification, grounding, and cabling
  - how full is the rack, i.e., how many RUs are in use and how many are available, if necessary, taking into consideration the room's load-bearing capacity.
- Electricity supply
  - current electricity supplies, their phases, fuses
  - locations of the power sockets
  - locations of the main distribution board and the fuse panels
  - loads of the phases
  - electrical load of the entire hardware room
  - contact information of the electricity provider and technical support
- UPS and reserve power
  - How are the uninterrupted power supply (UPS) and reserve power arranged?
  - How long will it take to have it available?
  - How long will the power supply work with the current electrical load powered by the UPS and reserve power supply?
    - If the reserve power is not sufficient for the required electrical load, servers need to be shut down. In that case, follow the “prioritisation list X”.
  - Contact information of the reserve power system's supplier and maintenance, and technical support
  - What kind of a testing arrangement does the system have, and who performs the regular tests?
    - Things to be tested: the operation of the reserve power machine's engine, starting automation, automatic switchover of the electricity supply, etc.
    - With regard to the reserve power supply machines, you must also monitor the amount and quality of fuel (fuel may deteriorate if it stands unused for a long time)
- Cooling
  - Where are the cooling machines located?
  - What is their cooling power?
  - How much cooling power is currently required?
  - Contact information of the cooling machine supplier and maintenance, and technical support
- Building technology automation, alarm systems, and monitoring systems
  - What kind of systems does the operation of the room depend on, e.g. cooling control systems, electricity supply automation?

    - What kinds of issues cause automatic alarms to be sent, and to whom?
    - Contact information of the parties/persons responsible for the above-mentioned systems
- Cabling
    - Locations, IDs, and connector types of the cable panels
    - Types, lengths, and routes of optical fibre cables
    - Categories, lengths, and routes of copper cables
- Access instructions
    - Access instructions for the hardware room
    - Personnel list including contact information, who is allowed access to the hardware room
    - Any other instructions related to access, e.g. key management.
- Security/fire extinguishing systems, and lighting
    - Document the information on the security systems
    - Instruct the users in the use of these systems during incidents

APPENDIX 2: Networking hardware, physical connections, address spaces

With regard to networking hardware, we recommend documenting the following things, for instance. This document does not discuss an expansion plan, construction of missing cable infrastructure, colour coding of the cables, or (hardware) locks.

- Device
  - Device type (router, switch, firewall, etc.)
  - Name (in the DNS system/nickname).
    - We recommend labelling the device with a uniquely identifying ID/name so that it can be quickly located when necessary
- Mark and model
- Capacity (e.g. 1G/10GE, can be updated to 100GE, etc.)
- Time of procurement/deployment (warranty period/service life estimate)
- Software version(s) and any software licenses
- Device owner/maintainer (group) (IT management etc., contact information, …)
- Physical location (hardware room, floor plan)
- Number of power supply units and how many are needed at a minimum for normal operation
- Maximum power requirement, thermal load
- Weight, size (L x D x H, RUs)
- Contact information for the device manufacturer / service, and instructions related to service/spare parts
- Connections
  - How is the device physically connected?
  - cable panel ID and cable pairs
  - Power socket ID and the electrical phase used
- Addresses
  - Address spaces and addressing for the devices and the local area networks
  - Local campus-specific VLAN
  - VLANs between several sites
- Management
  - Device's management address
  - Management program(s) (CLI, GUI)
  - User IDs and their access rights
  - Shell accounts and passwords or SSH keys
  - Allowed hosts/networks for management connections such as SSH, SNMP, syslog, etc.

APPENDIX 3: Basic network services

Document the basic network services that are absolutely necessary for basic telecommunications. Also, list how redundancy has been arranged for the services in case of incidents/maintenance. Where applicable, link the information to the device information documented in Appendix 2.

- DNS service
  - Server name
  - Physical location
  - Redundant server information
  - Hardware information incl. connections (cf. Appendix 2)
- DHCP
  - Server name
  - Physical location
  - Redundant server information
  - Hardware information incl. connections (cf. Appendix 2)
- NTP
  - Server name
  - Physical location
  - Redundant server information
  - Hardware information incl. connections (cf. Appendix 2)
- AD
- Secure-appliance
- Firewall
  - rules and their maintenance process
  - Hardware information incl. connections (cf. Appendix 2)

APPENDIX 4: Server hardware, disk servers

Document the licenses of the servers (and possibly their prices) the storage and backup/standby servers, server connections, and their interdependencies.

- Server hardware
  - Type of server (Blade, etc.)
  - Name (in the DNS system/nickname).
    - We recommend labelling the device with a uniquely identifying ID/name so that it can be quickly located when necessary
  - Mark and model
  - Capacity (1G/10G, can be updated to 100G, etc.)
  - Time of procurement/deployment (warranty period/service life estimate)
  - Software version(s) and any software licenses
  - Device owner/maintainer (group) (IT management etc., contact information, …)

- Physical location (hardware room, floor plan)
- Number of power supply units and how many are needed at a minimum for normal operation
- Maximum power requirement, thermal load
- Weight, size (L x D x H, RUs)
- Contact information for the device manufacturer / service, and instructions related to service/spare parts

- Connections
  - How and to where is the device physically connected
  - cable panel ID and cable pairs
  - Power socket ID and the electrical phase used
- Addresses
  - Device's management address

- Address spaces and addressing for the devices and the local area networks Local campus-specific VLAN
- VLANs between several sites

- Management
  - Device's management address
  - Management program(s) (CLI, GUI)
  - User IDs and their access rights
  - Shell accounts and passwords or SSH keys
  - Allowed hosts/networks for management connections such as SSH, SNMP, syslog, etc.

APPENDIX 5: Virtualisation environments

Document the setup and testing of the service (backups, dependencies, and redundant service instructions, shutting down the service, disaster recovery, licenses). Document the management and monitoring issues.

Abbreviations

AD	Active Directory
CMDB	Configuration Management Database
CPU	Central Processing Unit
DHCP	Dynamic Host Configuration Protocol
CLI	Command-line Interface
DNS	Domain Naming System
FitSM	Family of standard for lightweight IT service management
GE	Gigabit Ethernet
GUI	Graphical User Interface
IP	Internet Protocol
ITIL	Information Technology Infrastructure Library
NTP	Network Time Protocol
RU	Rack Unit
SNMP	Simple Network Management Protocol
SPOF	Single Point of Failure
SSH	Secure Shell
UPS	Uninterruptible Power Supply
VLAN	Virtual Local Area Network

References:

REF 1: http://www.fedsm.eu/fitsm
REF 2: https://en.wikipedia.org/wiki/ITIL