

WHITE PAPER

Distributed Printing with BS200/OSD and Dprint/Wprint/Xprint

Issue June 2009

Pages 9

Dprint belongs to the "Print Services" products of BS2000/OSD. It is used in the context of a network of servers and provides the following benefits:

- Optimization of the hardware and software utilization by sharing the printers among several BS2000/OSD systems and Unix systems
- Direct printing to or from remote sites of the enterprise or partners, with an easy interface, an automatic control of the resources and a complete follow-up and control of the print jobs
- Communication with Unix (via Xprint) and Windows (via Wprint) systems, allowing the above usage models in heterogeneous architectures
- High availability and reliability by rerouting print jobs and printers from a host to another one
- Seamless connection between the MS/Windows world and the BS2000/OSD mainframes



Contents

Highlights

DPRINT in homogeneous environment	2
Sharing the devices over the enterprise	3
Concentrating the print services on an unique Print Server	3
High availability with a work- and standby-server	4
Integrated document distribution with DPRINT	4
DPRINT in heterogeneous environment	5
DPRINT and Xprint	5
DPRINT as bsd-lpd Print Server	5
LAN to HOST printing	6
DPRINT and WPRINT	6
DPRINT Architecture	7
DPRINT Administration and interfaces	8
Cluster administration	8
Print Job creation	8
User interfaces	9

Highlights

Client/server architecture

DPRINT implements a consistently symmetrical client/server architecture within homogeneous (BS2000/OSD and BS2000/OSD) and heterogeneous networks (BS2000/OSD and UNIX systems). UNIX and Windows clients can submit print jobs to be processed on BS2000/OSD high-speed printers while BS2000/OSD can submit print jobs on printers attached to Windows via RSO.

Cluster model

A set of servers and printers connected via a LAN/WAN network are combined to form an organizational cluster. The major advantages of this model:

- **Single point of configuration:**
The cluster configuration is defined and modified centrally on the master host by the cluster administrator.
- **Single point of administration:**
All devices and jobs may be administrated centrally from the master host.
- **Single system image:**
Users do not even have to know how printers are distributed within the cluster. DPRINT handles automatically the server selection. In counterpart, adequate pool definitions allow the end users to address a specific printer or any printer controlled by a specific server.
- **Interoperability:**
DPRINT is able to send print jobs to remote systems, external to the local cluster like a remote BS2000/OSD cluster, a XPRINT domain. Symmetrically DPRINT can process jobs sent by such remote systems.

Network-wide printer access

DPRINT supports all spool subsystems on the server side (including SPS). The entire range of high-performance printers, such as Model2 (HP/HP90), Twin and APA printers, is thus available to the associated client hosts.

Cost benefits

Printers and resource sharing cut costs. Similarly, the print resources in a network do not need to be duplicated. A state-of-the-art architecture and standardized protocols ensure protection for your investment.

International standards

The print functions to be supported within heterogeneous networks are based on ISO standard 10175(-1) DPA (Document Printing Architecture).

Multilevel security functions – like privileged cluster administrator rights - and access controls for using the server mean that DPRINT satisfies all the stringent requirements of BS2000/OSD.

Reliability and availability

If a printer becomes unusable, DPRINT is able to reroute print jobs to other printers in neighboring data centers. So you can always maintain quality and keep your deadlines. DPRINT allows to transparently access printers that are switched between several hosts. The processing is thus independent of the availability of a particular one.

If the master host is down, DPRINT allows another host in the cluster to assume the role of master.

DPRINT in homogeneous environment

DPRINT provides the highest level of functionality in the frame of a network of BS2000/OSD systems. Schematically, all systems belonging to a site, on different hosts, or on different Virtual Machines on an host can constitute a DPRINT cluster with an integrated administration.

This is adapted to an enterprise with several BS2000/OSD servers, but also to an enterprise having a production and a development server.

For communication between distant sites, for which an integrated administration is not adequate, DPRINT provides the inter-cluster communication. This function justifies the existence of clusters containing only one system.

The main usage models are described in the following sections:

- Sharing the devices over the enterprise
- Concentrating the print services on an unique Print Server
- High availability with a work- and standby-server
- Integrated document distribution with DPRINT

Of course, the applications of DPRINT can merge several usage models.

Sharing the devices over the enterprise

When several BS2000/OSD systems are in use in the enterprise, DPRINT allows any printer of any type, controlled by one system, to be used by any application or end user located on another system. This allows the most appropriate printer to be chosen at any time. In addition, jobs or categories of jobs can be modified by a simple command, so as to reroute them to other printers.

The benefits are:

- Easy load balancing between systems
- The unavailability of a printer on a system has no consequences on the applications
- The enterprise can purchase a high capacity printer instead of several smaller ones
- The applications can run on the most appropriate system, independently of the connected devices
- An application can be moved from a system to another allowing a very high availability

In this example, the system running on host A has no printer. The DPRINT Client is installed allowing the print jobs to be transferred to the systems on host B and C which both run the DPRINT Server. On host B, one printer is "distributed". It is known by DPRINT and hence by any Client. The other one is exclusively used by applications running on host B. The same PRINT-DOCUMENT command is used for printing, independently of the print location. The choice is done through the Printer-Pool that contains local printers or distributed printers.

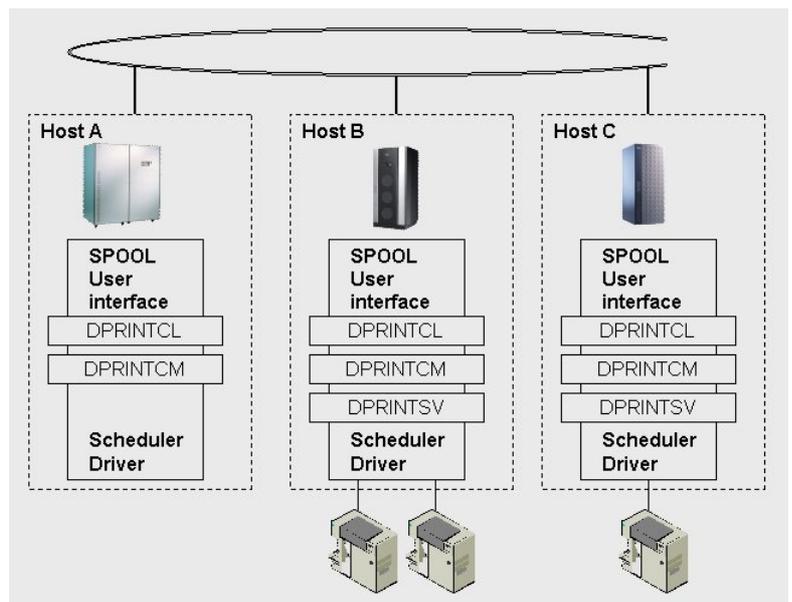


Figure 1: Printer sharing

Of course, printers belonging to a distributed pool can be connected to different hosts. DPRINT always searches the optimal printer: if a printer of the distributed pool is connected to the same host, and has the requested characteristics, it is selected by DPRINT. Otherwise, the job is transferred to another host with an adequate printer.

Remarks:

- A Spool option specifies the "default printer pool" of the PRINT-DOC command. This allows all spool jobs without explicit destination to be submitted to DPRINT.
- To select a print server, DPRINT also considers the job class of the print job and the Access Control Lists. This allows numerous tuning possibilities.

Concentrating the print services on an unique Print Server

This is a particular case of organization. The same host controls all printers. This host can also run other applications or it can be exclusively dedicated to the print services.

This kind of organization has also the advantage to concentrate all print activities of the enterprise in the same location with an unique point of control.

This is also the first step to the configuration described in the next section, where a pair of servers allows a very high availability.

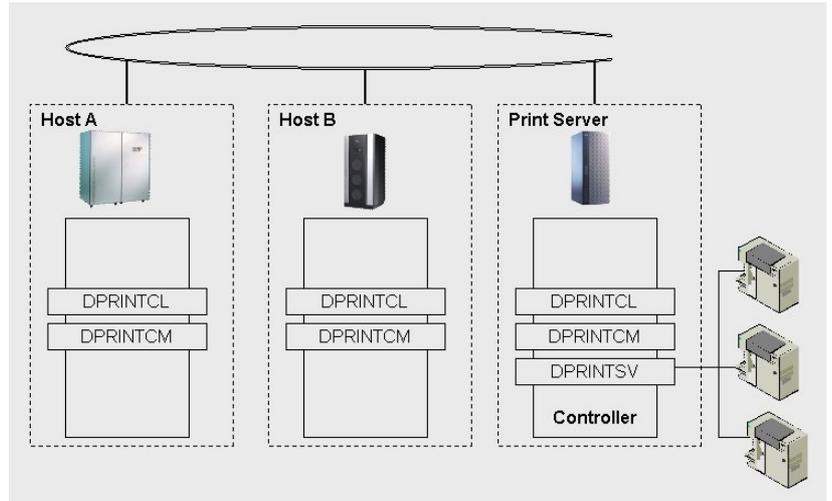


Figure 2: Central print server

High availability with a work- and standby-server

A printer can be switched from a host to another one. In usual configurations, this allows the printer to work successively for the different systems, and hence applications of those systems. DPRINT as described in the preceding section provides the same possibilities without the necessity of switching a printer between hosts. However the switch ability can be useful in the case of DPRINT to provide a very high print services availability.

Let's consider the following configuration: the printers can be switched to 2 different hosts, each one running the DPRINT server functionality. The physical printers have to be declared locally on each server, and declared twice in the DPRINT configuration (once for each print server system).

Of course, only one of the distributed printers is available at a time, since only one printer exists. There is however no problem in declaring both in the same distributed pool: DPRINT will never send a job to an host without active DPRINT server.

In this example, each printer is defined twice as distributed printer: once through Print Server X and once through Print Server Y. One or several distributed pools are available that contain both definitions of a given printer. The administrator starts the DPRINT server on the host that controls the printers. This does not prevent a DPRINT client to be started on the other host.

By this way, the end-user does not need to know which Print Server is active. DPRINT will choose the right one in all cases.

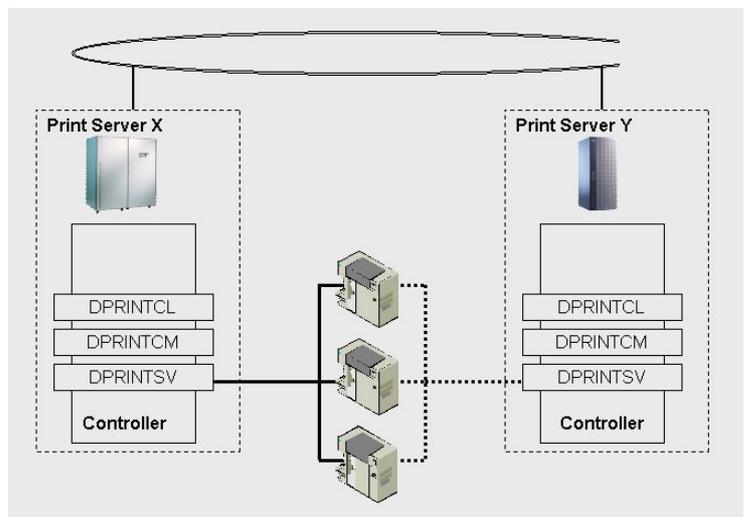


Figure 3: Work- and standby-server

What happens when the active server fails? In general, print jobs are already assigned to this server; some of them are even active. A DPRINT administrator intervention is necessary to transfer those jobs to the new currently active server. This is done by a single command allowing the transfer of the jobs satisfying a choice of selection parameters.

It is very important to notice that this operation does not need any information from the failed server. Until the job completion, the complete information is kept on the client.

Integrated document distribution with DPRINT

The use of DPRINT is not limited to a single site. DPRINT can connect systems arbitrarily distant. In this case, instead of considering a unique print server (as in the preceding examples), one considers several print servers - at least one on each site - and distributed pools allowing to directly identify the location of the server.

By this way, instead of printing 100 copies locally, and physically distributing the right number of copies to the distant sites, the user has just to give one PRINT-DOCUMENT for each site, using the appropriate printer pool name and copies number.

The advantages are:

- No physical transport
- Document electronically transported once to each site, and printed there the needed number of times.

Benefits:

- Sparing the cost for transport
- Delays only dependent on the network capacity; in the best case, the document is printed simultaneously in each site.

For this distribution, it is of course possible to build only one DPRINT cluster including all sites. This allows the administration to be done from a central point; this is probably better when the documents are all starting from the central point. This needs however a correct organization, in particular for the resources.

In other cases, the various sites rather need to be independent. In this case, DPRINT provides the Inter-Cluster communication. Each site is a separate DPRINT cluster (possibly with one single host). The cluster identification is to be specified in the PRINT-DOCUMENT command.

In this configuration, two clusters are communicating. In each cluster, one of the hosts has the role of gateway. Only the gateway network address is necessary to access an external cluster.

This address with an arbitrary cluster-name is to be entered in the parameter file. The cluster-name can then be used in the PRINT-DOCUMENT command.

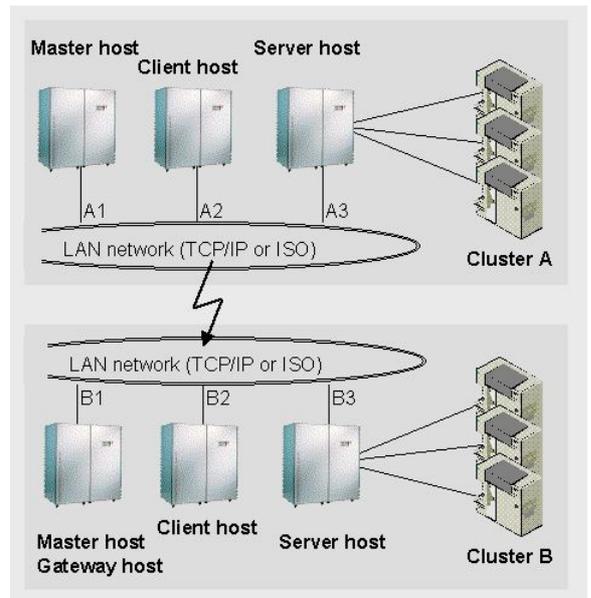


Figure 4: Distributed print servers

DPRINT in heterogeneous environment

DPRINT and Xprint

A DPRINT cluster is able to communicate with an Xprint domain in both directions. From a BS2000/OSD system, an Xprint domain is roughly comparable with an external DPRINT cluster.

The main differences are:

- the external cluster must be declared with the type "Unix", and its physical network address must be given
- in the PRINT-DOCUMENT command, an Xprint printer must be specified.

In the reverse direction, a Unix or NT user can submit jobs to the BS2000/OSD via Xprint.

Remark: the reference model selected as the basis for the DPRINT functionality was ISO DPA 10175.

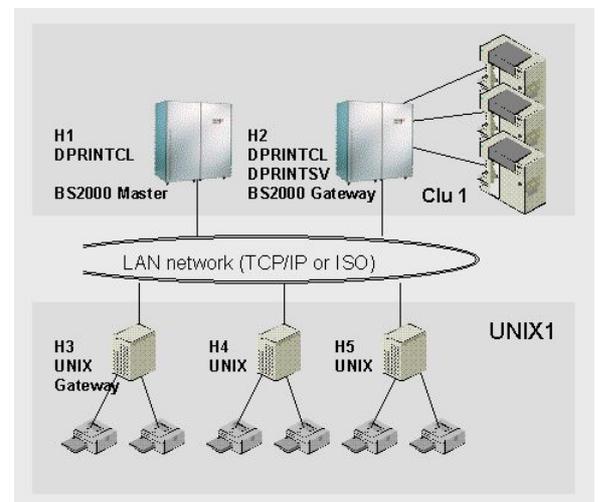


Figure 5: Dprint-Xprint coupling

DPRINT as bsd-lpd Print Server

DPRINT includes a bsd-lpd gateway. Any host supporting the corresponding client protocol is thus able to submit jobs to a DPRINT cluster, and control them.

In the lp job submission, the device name must be a distributed printer pool. If such a pool does not exist, the destination is supposed to be an RSO device or pool.

LAN to HOST printing

DPRINT and WPRINT

Wprint runs on Windows 95/98 and Windows NT and is able to submit print jobs to an external server via the bsd-lpd protocol. It is thus the ideal link between the MS/Windows world and the BS2000/OSD Print Services.

This feature is implemented in DPRINT and runs under POSIX. It consists of

- an FTS PageStream printer driver for Windows (integrated in Wprint or provided with DPRINT). It generates a device-independent data stream that will be handled later on by the gateway),
- a "BSD/LPD" gateway on BS2000/OSD that resubmits the Wprint jobs to the BS2000/OSD Print Services,
- a converter FTS PageStream / PCL, AFPDS or EXCCW.

All the RSO printers accessible from the Gateway, and all distributed printer pools are available from the Wprint client. This includes the AFP printers controlled by SPS.

Usage models

For the MS/Windows user, the Wprint/Dprint functionality is available in the "PRINT" interface of any product using the GDI interface. In particular, all MS/Office products, the Web browsers, Acrobat Reader, and so on.

In the printer names list, one or several items correspond to RSO printers or "Printer-Pool" name in a DPRINT cluster. The processing is the same as for using another printer.

A print job appears in the Windows Spooler interfaces, until the job has been transferred to the host. At this time, the job can be followed up in the Wprint interface (quite similar to the interface of the Windows Spooler).

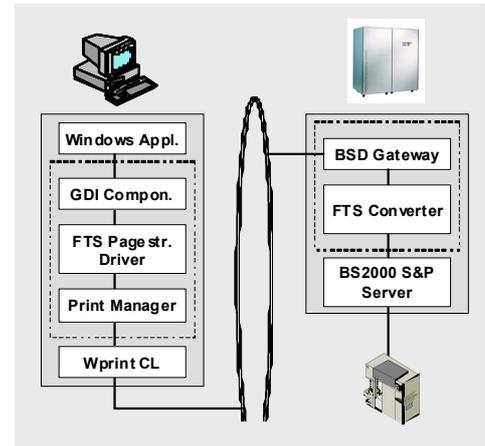


Figure 6: Dprint-Wprint coupling

Available drivers

The driver to be used on the PC can be

- an usual PCL or Postscript driver; the corresponding jobs must then be sent to corresponding RSO devices.
- a Generic Text Driver (MS/Windows component) that ignores any document layout or image. This driver allows printing in line mode on any printer controlled by BS2000/OSD.
- the FTS PageStream Driver that produces a device-independent data stream that can be converted into AFP, PCL or EXCCW. In this case, all aspects of document layout and images are taken into account.

The FTS PageStream driver provides the widest possibilities. In the frame of DPRINT, the conversion into AFP, PCL or EXCCW is done in accordance with the specified printer pool. In the frame of IDOM, the conversion is totally controlled by the IDOM application (see next section).

Customer benefits

- Wprint and DPRINT build a seamless link between the MS/Windows world and the mainframe. For example, MS/Office documents can be printed on production printers, as easily as on a printer directly attached to the PC.
- The FTS PageStream driver brings an additional flexibility.

DPRINT Architecture

DPRINT Cluster

DPRINT consists of a client subsystem (DPRINTCL), a server subsystem (DPRINTSV) and a common subsystem (DPRINTCM) that contains general functions. On each host of a cluster, the common and the client subsystem must be installed; the server must be installed on the hosts where print jobs are processed. All communications are done with CMX-TP.

Configuration file

A configuration file is present on each host of the cluster. It contains the description of the hosts, the printers, the printer pools, the access control list, and the remote clusters (BS2000/OSD or Xprint).

A BS2000/OSD host belongs to one and only one DPRINT cluster. The configuration file is thus unique on an host, with a fixed name. It is managed on the master host and distributed to each host of the cluster; this is ensured by DPRINT itself:

- At DPRINT startup time on the master host, the master host's configuration file is distributed to all other hosts of the cluster (where DPRINT is started).
- At DPRINT startup on a non-master host, the configuration file is imported from the master host (if DPRINT is started)
- during a session, all modifications performed on the master host are automatically distributed to all other hosts of the cluster.

Distributed printers and printer pools

DPRINT does not directly control printers; its role is to transport the information (files to be printed, parameters, possibly resources) from the client to the selected server and to submit the job to the local spool on this host.

A distributed printer is always linked to a local printer on a given host. DPRINT uses its own printer names to avoid potential naming collisions between the various hosts.

The distributed printers can be started or stopped on the master host by the cluster administrator or on the host where it is connected by the spool administrator.

Distributed printer pools gather distributed printers that may be linked to different servers.

Creation of a DPRINT cluster

To create a cluster, DPRINTCM must be started on all hosts intended to be included in the cluster. A "master" host is to be selected, on which all hosts of the cluster are to be declared by specifying their BCAM name. This is done with SPSEVE. On each other host, the "master" host and the local host must be declared by specifying their BCAM name. This is done with SPSEVE. Furthermore, access rights must be set to *ALL for each host in order to allow communication of one host to another.

Once this is done, the cluster creation is launched by the "master" host. Later on, any configuration and administration are done from the master. Every action is automatically replicated on all other hosts.

The "master" is unique at any time, but if the present master fails, the master function is automatically taken over by another one where the configuration and administration can be continued.

The master can also transfer its privilege to another host by command. The current host loses its privilege, while the new one becomes the master of the cluster.

Configuration

One or several hosts of the cluster can share some of their printers across the cluster. For each of these hosts, the cluster master has to define a server in the configuration. Each server is described by its logical name, the host name it belongs to and various access rights. Then, the master has to define the shared printers named distributed printers associated to the different servers. The parameters defining the distributed printers are the server that controls the printer, the local printer name and its distributed printer name. Distributed printer pools are defined by specifying the pool name and the names of the distributed printers that belong to the pool. The master may also declare external clusters that can be reached from his local cluster. For that purpose, he has to define all these ones as DPRINT Remote Clusters by specifying a logical name that will be used for later access and an address composed of a BCAM name and an internet address. Those remote clusters can be BS2000/OSD clusters or UNIX clusters.

All the definitions are performed with SPSEVE.

Accessibility

In order to control the accessibility of the cluster elements, the master can define specific access control list objects. These access control lists can be associated to the hosts, the servers and the remote clusters of the configuration.

All the definitions are performed with SPSEVE.

DPRINT Administration and interfaces

Cluster administration

During a session, the administration of the cluster is done with SPSEVE.

A set of statements is available to add, modify, remove the various objects belonging to the configuration, that are:

- the remote cluster
- the hosts
- the servers
- the printers
- the printer pools
- the access control lists

The cluster configuration (add / remove hosts from the cluster, changing the master, ...) may be done by specific DPRINT commands.

Show commands and statements are also available. Some of them are also available to the end user.

Print Job creation

The jobs are created by the PRINT-DOCUMENT command of the BS2000/OSD Spool.

Depending on the current default spool options, the following parameters must be considered:

- in a single system image configuration, no particular parameter is to be used; the job may be submitted to DPRINT in all cases provided that it has been specified in the global spool parameters panel.
- in the other cases, the locality of the print job depends on the parameter-values given in the "TO-PRINTER" parameter:

Cluster-Name	Printer-Name	Remark
*LOCAL-CLUSTER	*STD	The job is submitted to the default destination defined in the global spool parameters panel. See next table.
*LOCAL-CLUSTER	<RSO printer> or <RSO pool name>	The job is submitted to the specified RSO printer or pool
*LOCAL-CLUSTER	<local pool name>	The job is submitted to the local printer pool
*LOCAL-CLUSTER	<distributed pool name>	The job is submitted to the distributed pool
<BS2 cluster name>	<distributed pool name>	The job is sent to the external cluster. It is submitted to the distributed pool.
<Xprint cluster name>	<Unix printer name>	The job is sent to the specified Xprint domain. It is submitted to the specified printer.

Global Spool parameters ruling the default print job destination:

Dprint-Destination		This parameter defines the default destination for any print job for which no specific PRINTER-NAME parameter has been specified. It defines if DPRINT has the priority or not
*CENTRAL		All the print jobs are addressed by default to DPRINT cluster considered as a single system (Priority is given to DPRINT)
*POOL(POOL-NAME=)	<distributed pool name>	All the print jobs are addressed to the DPRINT printers pool specified as pool name (Priority is given to Dprint)
*DESTINATION		The priority is given to the local Spool or RSO (To address a specific DPRINT server, it is necessary to specify a DPRINT printer pool explicitly in the print command)
Destination		
*CENTRAL		All the print jobs are addressed by default to the local Spool
*REMOTE(DEVICE-NAME=)	<RSO printer name> or <RSO pool name>	All the print jobs are addressed by default to the specified RSO printer or printer pool
*POOL(POOL-NAME=)	<local pool name>	All the print jobs are addressed by default to the specified local pool name

User interfaces

For the end-user, the usual SPOOL commands are to be used to manage the distributed print jobs.

The PRINT-DOC command allows to create a distributed print job which receives a local job-id that is to be used for all actions related to this job (SHOW, MODIFY, CANCEL).

In general, the commands concerning a list of jobs contain selection parameters allowing to specify that the action concerns the distributed jobs or does not concern them or concerns certain distributed jobs, and so on.

It is important to remark that the "MODIFY-" command has the effect to re-evaluate the best host to process the job, even if no parameter (except the job identification) is given in the command.

For the end user, this command can only be executed when the server is available. However, an interesting feature allows the cluster administrator to redirect a print job when a job has been sent to a server that fails or is shut down before the execution of the job. So, this simple command has the effect to redirect the job to a server available at this time.

A MODIFY command exists that allows redirecting a set or all the jobs sent to such a server onto another available one.