
What is DSSUltra?

DSSUltra is a multifaceted set of enhancements and extensions to z/OS DFSMSdss (hereinafter referred to as DSS) data backup services, providing increased automation, performance, and functionality, while reducing administrative intervention and overhead. DSSUltra does not replace DSS, but rather augments its existing capabilities. DSSUltra features include automated recognition of the entire DASD backup inventory, automated retry of failing backup operations, automated support of z/OS and non-z/OS (e.g. z/VM and z/Linux) volume formats, backup data FTP offsite for disaster and other data recovery purposes, and extended parallelism of backup and FTP processes. DSSUltra implements its features in the context of DSS backup to DASD. In installations where DSS is being utilized to perform backup to tape, the benefits which the DSSUltra implementation offers can be demonstrated to be superior to those associated with tape backup. Alternatively, DSSUltra can operate in conjunction with existing backup to tape to realize its additional exclusive advantages.

Throughout the remainder of this summary, the terms “backup” and “dump” are used interchangeably.

DSSUltra is a product of Expanded Systems Services (ESS) Ltd. in Calgary, AB, Canada.



DSSUltra Benefits

DSSUltra offers the following enhancements and extensions to DSS:

- DSSUltra is constructed to, by default, recognize and process all online DASD volumes without the need for explicit selection specifications. The advantage of this approach is that volumes can be added, changed, or deleted; and those added or changed will be, or will continue to be, automatically backed up without requiring changes to the backup jobstream. However, should selection specifications be required, they are supported generically via wildcarding, as well as both inclusively and exclusively. For example, IPIV=ESS- will include only volsers beginning with ESS in the backup, and IPEV=ESS0- will exclude the subset of IPIV=ESS- volumes beginning with ESS0. Additionally, DSSUltra supports the backup of one or more SMS storage groups via IPIS control statements. Any combination of IPIS, IPIV, and IPEV specifications may appear in the control statement stream, providing maximum explicit selection flexibility when required.
- When output to nonSMS volumes is chosen, wildcarded selection specifications of the volsers to which the output backup datasets are to be directed are supported. For example, OPIV=BU*0- will result in output backup datasets being directed to volsers whose first two characters are BU, whose third character is any valid in a volser, whose fourth character is 0, and whose remaining characters are any valid in a volser. When SMS output is chosen, all volumes in the storage class specified on the OPIS control statement are eligible to receive output backup datasets.
- DSSUltra employs a flexible automated GDG-based mechanism for creating the backup datasets which are used during the course of its processing. This includes automated definition of the GDG bases themselves, obviating the requirement for separate IDCAMS executions which would otherwise be necessary to perform such definitions.
- Upon completion of the first full backup of each volume, DSSUltra captures and retains the volume's status by creating a copy of its VTOC, termed a historical VTOC copy. Prior to initiating its next full backup, DSSUltra, if requested, compares the historical VTOC copy with the currently active VTOC on

the volume. If the two are identical, the backup is bypassed, saving time and storage. If not, the backup is performed, and upon its completion the historical VTOC copy is updated from the just-backed-up VTOC in preparation for the next cycle. This is, in essence, an “incrementalized full volume” backup.

- DSS implements multiple parallel backup processes through specification of the PARALLEL control statement, which results in the creation of a new subtask for each individual backup operation. The created subtasks execute in the DSS jobstep address space or, if the cross-memory application interface is requested, in a separate server address space which is launched in response to the request. DSS does not, however, launch multiple server address spaces, i.e. it does not schedule each individual backup operation into a separate address space. While such an approach incurs some additional overhead, it offers benefits in the form of increased reliability as a result of individual process isolation, and increased control through the ability to more easily manipulate individual processes executing within address spaces rather than subtasks. DSSUltra utilizes the multiple address space approach.
- As noted above, DSSUltra can back up both SMS and non-SMS volumes, with the output backup datasets being directed to either SMS or non-SMS volumes. In both SMS and non-SMS output scenarios, DSSUltra creates multivolume datasets in an attempt to utilize DASD space as optimally as possible. However, when non-SMS volumes are selected for output, the multivolume datasets which DSSUltra creates can be tailored more precisely to match available output space. This results in a potentially significant improvement in space utilization efficiency.
- DSS itself is capable of backing up volumes in formats other than z/OS, e.g. those in z/VM or z/Linux format. However, without DSSUltra, this capability is not dynamic. DSS must be informed prior to execution of the backup jobstream, via an option in the associated backup control statement, of a z/VM-format volume. If DSS has not been so informed, the backup will fail or be incomplete. DSSUltra analyzes the VTOC of each volume prior to dumping to determine if it represents a z/OS, z/VM, or z/Linux volume, and adjusts the dump options accordingly. In addition, DSSUltra reclaims control at the conclusion of each dump operation to assess its success. In the event of an error, it will retry the operation with alternative options in order to maximize the probability of obtaining a valid and usable dump.
- DSS does not list the names of individual datasets being dumped in the course of a DUMP FULL operation. DSSUltra addresses this omission by capturing and listing the individual dataset names.
- ESS strongly recommends use of the DSS CONCURRENT keyword to maximize backup data integrity, and to avoid the DASD RESERVE which DSS issues in the absence of a CONCURRENT specification. However, in emulated DASD environments such as Hercules or FLEX-ES, where the emulated storage subsystem does not support CONCURRENT, the RESERVE is still issued. DSSUltra offers an option to disable the RESERVE and avoid the inevitable associated device lockouts. While this approach compromises backup data integrity, the option may still be desirable in situations where “fuzzy” backup data is deemed acceptable.
- DSSUltra can launch up to two independent concurrent subtasks which transmit output backup datasets via standard FTP to up to two z/OS, Windows, Unix, or other FTP-supporting destinations. These destinations can include removable external hard drives attached to Windows or Unix workstations or laptops. The transmissions occur in parallel with backup operations, and facilitate the creation of offsite copies of backup datasets for disaster or other data recovery purposes. Additional flexibility permits wildcarded inclusive or exclusive specifications of the volsers which are to be FTPed, if subsets of the full spectrum of backed-up volsers are required. DSSUltra also offers the option to delete output backup datasets from DASD after successful FTPs, if disk space reclamation is desired.