

A DATA CENTER WITHOUT ENFORCEABLE STANDARDS RISKS MUCH MORE THAN MERE STORAGE MISMANAGEMENT.

Efficiency and reliability are the twin pillars of mainframe computing. Here's why you need storage allocation control standards, and four things to help you keep tabs on them.

On December 14, 2020, users of Gmail, YouTube, Google Drive, Google Maps, Google Calendar, and more — in other words, most of the online world — couldn't access these and other services for around 45 minutes. The problem, as the search-engine-turned-tech-titan later explained, was an identity management system that [ran out of storage space](#) and lost the ability to authenticate user logins. Compounding the issue was the fact that it affected Google's internal employees as well, who unsurprisingly rely on tools such as Gmail and Google Meet to manage works in progress and service requests. To remediate the issue, or at least temporarily work around it, teams disabled an automated storage quota management system that had failed to properly allocate the necessary space for identity management demands. Google's distributed environment means the company has never utilized mainframes, but the outage underscores the importance of proper storage monitoring and management regardless of computing platform. This could easily have been a mainframe use case.

Storage is critical for every single function of an application. Without storage, applications don't exist. If storage isn't available or an application runs into issues trying to access it, that app will, at best, experience throttled performance, or at worst ABEND. Complications with storage can also lead to data unavailability, which can even become permanent data loss in severe circumstances. Clearly, storage is one of the primary pillars of application integrity.

While your own storage management mishaps might not be felt around the world, they'll no doubt be a major inconvenience for the consumers of your business services, and possibly an even bigger pain for you. To avoid service disruptions and production mishaps, it's important to maintain strict standards governing how the different datasets in your installation are treated throughout their lifecycle, from creation and processing all the way through to deletion. That's where DTS Software's Allocation Control Center, or ACC Monarch, comes in.

ACC Monarch provides centralized, rule-based management across a wide variety of mainframe resources, enabling you to not just establish standards, but to enforce them throughout the dataset lifecycle in both test and production environments. With ACC Monarch, you can define the attributes of datasets at their creation, manage access and buffering as they're used, and even dictate how they're disposed of according to your production and security requirements.

Storage allocation standards enforcement might not sound all that exciting, but storage management integrity has a direct impact on application reliability and performance. On the mainframe, where the world's most mission-critical business services are delivered, standards enforcement is vital to these services that we consume every day. ACC Monarch can be used in conjunction with your allocation policies as a central repository for the standards you require throughout the storage management process.

GAIN GREATER CONTROL OF YOUR DATA CENTER WITH BETTER STANDARDS ENFORCEMENT

It's not uncommon for a busy data center manager to settle for suboptimal performance in an area because it's considered too time-consuming and expensive to make changes to JCL, control statements, or the DFSMS environment in the process of improving a job that already runs in production. It is a fair assessment when any change to the production environment requires a lengthy rollout through development and testing phases. But what if there were an easier way?

It's possible your JCL may have been written 30-40 years ago under a completely different set of standards (*and their enforcement methods!*), but that doesn't make it any less

of an organizational asset. Standardizing JCL, eliminating defects, and improving quality will help application performance and ensure you get through thousands of batch jobs error-free, day in and day out. Instead of having to directly alter JCL so that it adheres to certain standards — a process that sometimes feels like it requires a court order — ACC Monarch works by executing on a set of policy rules. These policy rules control how datasets are used, enabling you to adjust the following features of your data center:

1. LEVERAGE I/O PERFORMANCE AS A BENCHMARK

Besides CPU consumption, I/O performance is the most important benchmark in data center management, because job streams can stack up and fail without fast and efficient I/O. In a typical IBM® z/OS® installation, sequential batch jobs and reading and writing VSAM datasets consume significant I/O resources, and it's up to data center managers to ensure that these resources are allocated effectively. The reality, however, is that many jobs, steps, and programs in a typical data center suffer from poor performance because they rely on requirements that might have been appropriate in 1975 but are obsolete today.

It's understandable that data center managers in the risk-averse mainframe world might not want to make what they perceive to be invasive changes to JCL or control statements when production jobs are able to run fine as-is. Fortunately, ACC Monarch's Performance Buffering feature can intervene to automatically increase buffering values for both non-VSAM and VSAM datasets without necessitating any JCL changes. In the

case of non-VSAM datasets, a simple increase in values like BUFNO can offer a substantial performance improvement, while system-managed buffering values including ACCBIAS and SMBVSB can control buffering across all types of VSAM I/O. By improving buffering performance, data center managers can reduce elapsed time and CPU time, increasing operational efficiency and ultimately saving their organizations time and money.

2. DATASET CREATION – NAMING CONVENTIONS STANDARDS ENFORCEMENT

Naming conventions are the first means that data center managers have to ensure the right data shows up at the right time to fulfill a specific function. Logical and consistent naming conventions are vital to efficient and reliable data center operation, which is why ACC Monarch Unit Name Manager enforces naming standards from the moment datasets are created.

ACC Monarch Unit Name Manager can validate or override UNIT name, but it can also examine and validate allocation variables such as volume serial number, DD name, dataset name, Data Control Block (DCB) and space characteristics, including RLSE. If necessary, the values associated with these variables can be dynamically changed to meet system requirements, ensuring the validity of not just UNIT names, but all aspects of dataset allocation and job execution.

3. DATASET DELETION IS NEVER THE END

Any time a dataset is deleted, there are options as to how the dataset is to be treated. For example, coding DELETE simply removes the pointer to the dataset, but coding the ERASE parameter writes 0s over

everything contained in the dataset, which can cause major CPU and I/O performance issues for large VSAM files. When a developer at a financial firm accidentally coded ERASE for a large dataset, it created a backlog behind this twenty-minute job that caused critical and expensive delays in other production jobs. ACC Monarch helps prevent this problem by enforcing standards that either limit or require operands such as ERASE, DELETE, and PURGE for disposal of various types of datasets.

In the case of VSAM datasets, the IDCAMS_DELETE environment gets control when VSAM clusters and non-VSAM datasets are about to be deleted. Standards enforcement rules might look at the job or userid that initiated the delete to ensure they have the right administrative privileges, or they might look at the characteristics of the dataset in question to determine whether or not an ERASE would cause CPU and I/O performance issues or disruption to critical production workloads. These measures ensure that a single mistake can't have consequences that are felt throughout an entire organization.

4. GDG PROCESSING

To ACC Monarch, generation data groups (GDG bases) and generation datasets (GDSS) are just another dataset, meaning ACC rules can provide values for primary and secondary space, dataset organization, DCB characteristics, and pool, storage group, or volume serial number. ACC policy rules can set characteristics of GDG bases in the DEFINE_GDG environment, which establishes control whenever a GDG base is defined or altered. These rules can enforce standards dictating the number of generations, roll off options, first-in first-out (FIFO) or last-in first-out (LIFO) options, and more.

Since Extended GDGs (GDGEs) were created in z/OS 2.2, released in September 2015, one of the most advantageous uses of ACC Monarch's DEFINE_GDG environment has been to enforce the creation of all new GDGs as GDGEs, ensuring that they can be defined with up to 999 generations, instead of the 255 you can reference with the original GDGs. This means that all of the existing jobs will use the new standard automatically instead of requiring someone to manually reach in, find the part of an application that creates GDGEs and make the necessary alterations. GDGE is far from mandatory — after all, you've lived without it since the 1960s — but the ability to easily make it the new standard offers you some additional features and use cases that you've been missing out on. Installations with a mix of old and extended GDGs with a variety of attributes are also more prone to error and likely to suffer from poor performance, which means standardization is advisable.

THE POWER OF POLICY ENFORCEMENT RULES

ACC Monarch relies on the DTS Software Policy Rules Engine, which allows it to enforce systemwide standards in such a way that they can't be bypassed under any circumstances. These rules can run before and after DFSMS ACS routines, as well as at many other points in the dataset life cycle, such as volume selection,

obtaining new extents and volumes, deletion, and cataloging. And ACC Monarch can record all actions in logs and SMF records. You can use ACC Monarch to issue WTOs and console commands, send emails, and more. Also, startup, refresh, and shutdown are all managed via simple commands and dynamically installed code, removing the need for SMP/E work.

ACC Monarch maintains control at every point in a dataset's lifecycle, from OPEN and allocation to extend, rename, and DELETE. Do you want to prevent the use of production resources like RLS log streams in test environments? No problem. Would you like to increase and standardize I/O buffer usage at OPEN? Consider it done.

These use cases are just the tip of the iceberg for ACC Monarch, and we're confident that it can help you manage your data center more effectively, no matter your goals.

FOR MORE INFORMATION

about how ACC Monarch can improve the efficiency and reliability of your production workloads through effective standards enforcement, [reach out for a demo](#) or to request a free trial at DTSSoftware.com.



4350 Lassiter Ave. at North Hills, Suite 230, Raleigh, NC 27609
www.DTSSoftware.com • info@DTSSoftware.com
918-833-8426 • 919-833-2848 (fax)