



FIM+ CPU Resource Utilization

Synopsis: The impact of FIM+ CPU usage at most sites should be negligible. This is borne out in existing production customers and is a result of FIM+ having performance features built in since its inception.

These features include:

- Offload hashing calculations to crypto cards found on all mainframes.
- Full scans run at low usage time. Quick scans run during peak usage.
- Control of when scanning is done resides with the customer.

FIM+ is highly efficiency at all times. The CPU time consumed should have no impact on capacity plans.

FIM+ CPU Time Usage: A quick scan reads only the meta data portion of the file. A full scan reads each block on the target dataset and requires more I/O and associated CPU time as a result. It is recommended that full scans are run occasionally (perhaps weekly), in off-peak times. Quick scans can be run many times daily due to ultra-low utilization. Scans can be initiated on-demand or scheduled at customer selected times.

In both types of scans, hash instructions are offloaded to crypto cards on IBM's z processors and do not accrue traditional CPU time.

CPU Example: The image shows the CPU time and amount of data read for a scan of the entire z/OS APF list for a standard configuration. Line 1 shows the CPU for a full scan of the same information and requires 17 seconds of CPU time. As expected, the I/O required is also much lower with a quick scan.

EntCt	Size	Cpu
35653	1040M	16.961
35653	452K	0.481

Line 2 shows a quick scan of the 35,653 members and takes less than ½ second of CPU running on a z16 processor.

Some customers run weekly full scans with ease and thousands of quick scans for critical applications during normal daily operation.

FIM+ can also scan and verify backup and log files which are sequential datasets. These types of scans typically run at approximately 10 CPU seconds per gigabyte of data stored. For very large backups, image copies and other log files we recommend our much faster sampling scans or hyper-speed integrity verify processing.

Memory usage: In general, FIM+ will manage the workload to fit the amount of memory allocated that you give it. Most people specify REGION=0M in the started task or batch job for both the FIM+ Server and the FIM+ Agents. The minimum requirement for the FIM+ Server is 64 Megs and for the FIM+ Agent 32 Megs and this will support scans containing about 100,000 members. If you are running many scans concurrently, FIM+ will delay scans automatically if available memory becomes an issue.

DASD Usage: As a rule of thumb, FIM+ stores about 200 bytes of data per file or member being scanned. This data is stored in an encrypted PDSE called the FIM+ Trust Vault. If the vault becomes full, FIM+ provides automated storage management that will allocate more space and resume operation.

As always, results in your site may vary and should be validated during a trial.