

# Platune

## Take Back the Data Center

Kelly L. Vogt

It was the worst ever. Month-end was here and they had laid siege to us. Egad! Now I know how they felt at Dunkirk. We fell back to the machine room. They charged, lobbing jobs at us at horrific rates. "Who is it," cried management? "It's the *ad hoc* users," we replied, "and a cohort of developers on our flank!" Management turned and ran. "O.K. men," I rallied. "It's up to us now. We're going to take back the data center!" -- This paper details how a squad of performance management specialists organized and streamlined test batch operations using automation and other vendor products.

### ...Overrun

"You guys hold this position," I shouted over the crushing din. "I'm going out to reconnoiter." I dashed out of the machine room, where we had holed up, and made my way to our office area. A phone went off near me; instinctively I hit the deck. I crawled over to it and checked the caller ID. Sacred bovine! It was Bob Zimway, one of our *ad hoc* power users. I gritted my teeth and lifted the receiver off the hook like I was disarming a booby trap. "Hello," I croaked into the phone. He instantaneously commenced his verbal incendiary bombardment. I did my best to take the call, cramming my words in between his expletives. "Yes, Mr. Zimway, I know who you are," I grated. He went off again like a phosphorus grenade. "We're doing the best we can," I rasped. He was having none of it. "A group of actuaries launched a surprise attack early this morning and submitted a whole raft of long-running jobs." I was shouting by now. "There was nothing we could do." I heard him fill his lungs deeply with a witty rejoinder, so I cut him off. "We'll review the situation and do what we can" I said, and quickly hung up the phone as hard as I dared.

I checked SDSF and TMON before I dodged my way back to the machine room to report. I signaled a huddle and my fellow squad members rounded on my position. There was Sergeant Vincent, management liaison and group leader. Jurgenson, our electronic surveillance specialist, qualified on all systems monitors. Meyer, counter-intelligence expert, he worked undercover in development back in '88; he knew the deepest recesses of their minds. And my buddy, Brewer, he defected from open-systems a few

years ago. Fortunately he got out in time, before any real damage had been done. I was on point that day.

"It's bad," I spat out, gasping for air. "The queues are backed up with jobs from all departments. It looks like most of the initiators are tied up with long-running jobs the actuaries submitted this morning. If any more jobs like that initiate... throughput will go to zero." We gave each other hard looks as I caught my breath. "There's more," I said in a warning tone. They studied me intently. "They've got the tape drives. They got all of them! They never had a chance..."

These were seasoned veterans, but this was hard to take. "Are there allocation waits," Sergeant Vincent asked dourly? I couldn't look him in the eye. "Yes," I said quietly to the raised floor. No one spoke for a few minutes. You could cut the tension with a hacksaw. When young Brewer suddenly broke the silence, his words tore the air. "We can't let them do this," he barked. "No, we can't," agreed Vincent. Jurgenson and Meyer perked up. "It's time we gird our loins for battle," Vincent ordered with savage determination. Suddenly, the room became bright with hope. "I have some ideas I want to work on," I said. "I'll check TMON and see what I can do to get some more cycles back," said Jurgenson. "And I'll make a few phone calls," added Meyer. With shoulders back and chins up we strode out of the machine room and returned to our burlap bunkers, filled with a sense of purpose. We had a mission.

### Regroup...

Sergeant Vincent ordered our squad out on patrol and we discovered a number of problems. Several

opportunistic actuaries submitted large numbers of jobs at dawn thereby monopolizing the initiators. Meyer's intelligence report indicated they entered the office under cover of darkness. It was hard not to admire such a dedicated adversary. Many of these jobs were long running. Once a long-running job started, it rendered the initiator ineffective for processing other jobs. One by one, initiators fell to occupation forces. Jobs piled up behind them like refugees at a border crossing. We were hard-pressed to even react to this dastardly tactic. Our only option was an end-run around WLM, manually starting jobs with the \$SJ command. This is an effective counter-measure, but you can't use it too often. We run WLM-controlled initiators allowing WLM to optimize initiator management.

Initial media reports indicated that the actuaries had taken control of all tape resources, but these reports were found to be gross speculations, foisted upon us by embedded news correspondents. Brewer infiltrated the tape library to serve as forward observer and discovered that developers were also conducting extensive data recovery operations. These jobs captured and controlled numerous tape units. Brewer also noted that the local citizenry were conscripted to fetch countless tapes one at a time, dooming them to drudges. He could see no way to free them without a full, frontal assault that was sure to result in massive job cancellations.

Jurgenson made extensive use of his TMON electronic surveillance equipment and discovered that the initial, swift progress of many jobs had been summarily halted. His full-bore drill-down indicated significant data set contention between developer jobs. He made generous use of his RMF Monitor III to confirm his sighting. He also reported that hostilities had spilled over into the production area, causing delays in critical production batch processing thus destabilizing the whole region. Jobs in the production area are plainly non-combatants. What do we have to do, rename them all UN\* something? Don't these barbaric developers have any respect for accepted codes of conduct? Our Shop Standards Manual, regulation UR2-L8 clearly states: "No non-production job shall allocate any production data set during such time that said production data set may be allocated for exclusive use by aforementioned production job." There would be no avoiding casualties this day.

### **War Plans...**

Once all bogies were identified and isolated, strategic command formed a task force to study the problems and devise plans of attack. Our initial thrust would be to out-flank the actuaries and regain control of the initiators. Once this ground was secured we would advance on the tape library, scatter the developers

and liberate the local inhabitants. Our final push would be to secure and protect our production data sets. This was critical to the overall success of our mission. Supreme command issued strict orders that production data sets would be protected at all costs.

We recognized that the enemy had deeply infiltrated the company. It could be anyone: the polite fellow at the water cooler; the nice lady who worked down the hall. Friend or foe? It was impossible to tell them apart just by looking at them. We needed intelligence and we needed it fast. We turned to our CA-NeuMICS database for information, where we kept detailed dossiers on every user and every job in the system. We began mining data to separate the decent citizens from the irregular militia from the subversive groups. We were hunting for jobs of mass consumption.

Examining history data, we learned the routines of our users. We categorized their jobs by resource requirements and arrival rate. We classified jobs by CPU time. We grouped jobs by tape requirements. We knew more about their habits than they themselves knew. Given all this information, we rendered our target coordinates. We knew just where to strike and how. We aimed our Photonic Cannon and began to fire!

### **Counter-Attack!**

Actually, we don't have a Photonic Cannon. But we do have a big gun in our arsenal. It's our MOAB; our Mother of All Batch tuning tools. We have a product called ThruPut Manager.

We launched our first counter-strike using ThruPut Manager Job Limiting Services (JLS), with a scheme to control the number of concurrent initiators a user can occupy. JLS provides a tool called a job-limiting *agent*. A JLS agent can be used to represent a system resource. A JLS agent has a threshold value associated with it. One or more JLS agents can be tagged to a job. When a JLS agent is tagged to a job, a weight is assigned, which represents that job's use of the agent. JLS will only allow a job to initiate when the job's added weight does not cause the agent's threshold to be exceeded. For our initiator management strategy, each job was assigned a JLS agent like INITS.userid with a weight of one and a conservative threshold of three. This effectively capped the number of jobs a user could run concurrently to three.

Thus far our operation proved quite successful. A few minor political skirmishes broke out, but they were quickly extinguished. Most people agreed that three concurrent jobs were fair and generally adequate. We received some complaints that once three long-running jobs initiated, the user was unable to do any

other work – even a quick print job. This seemed a reasonable complaint, so we added another JLS agent called LARGEJOB.userid with a threshold of two. This agent was tagged to a job depending upon the job's CPU requirements. Fast jobs<sup>1</sup> did not receive this agent; all other jobs were tagged with it. The LARGEJOB.userid agent ensured that a user's short jobs would always be processed in a timely manner. After this scheme burned in, we were able to increase the INITS.userid and LARGEJOB.userid thresholds to five and four, respectively.

### ***Secure the perimeter...***

Sometime later, Meyer tickled his covert assets and discovered a developer plot to probe our perimeter for weaknesses. He captured and interrogated a programmer trainee who revealed that developers were running clandestine jobs under alternate user IDs by way of the CA-7/Personal Scheduling product. These people love anarchy like they love to breathe. Their actions propagandized the job queues making target identification difficult for the Scheduling department, which is responsible for tracking these queues. In response, we took advantage of a field in our ACF2 user LIDREC<sup>2</sup>. In our shop, jobs submitted from the CA-7/Personal Scheduling product all have similar job names and user IDs like UCC####X. The UCC#### LIDREC contains a locally defined field where we store an alternate user ID. This alternate user ID matches the person's normal working user ID. We laid down a smoke screen and, using mirrors, coded ThruPut Manager exit 1 to extract this field and store it in a ThruPut Manager variable. Now UCC jobs are tagged with the INITS.userid agent using the alternate user ID, which closed the bulge in our front lines.

After a brief furlough, we mounted our invasion of the tape library. Our earlier intelligence work and experience with initiator management revealed that our JLS agent maneuver would be extremely effective in the close combat that is tape management. We devised a tactic to tag jobs with JLS agents like MAX3490.userid, MAX3590.userid and MAXVIRT.userid, representing our 3490, 3590 and virtual tape pools, respectively. The agent weights were set to the maximum number of tape units of each type used in a job step. Thresholds for these agents were set to allow the user a maximum of eight concurrent 3480 tape units, four concurrent 3590 tape units and twelve concurrent virtual tape units. We made a leaflet drop announcing the new standard, then fired for effect. We fought from silo to silo, but we

---

<sup>1</sup> Jobs, that require 13 CPU seconds or less. See my CMG 2002 paper, How We Successfully Converted to WLM-Controlled Initiators.

<sup>2</sup> Logon ID record; the anchor point for each defined user.

didn't lose a single drive unit. Some occasional tweaking of agent thresholds is required when belligerents request additional resources for special-forces projects. This is a simple matter of issuing console commands, when necessary.

Having acquired command and control of the tape hardware, we now focused our efforts on liberating the tape librarians. Taking this objective would require overwhelming force, heavy ordnance... a bunker buster. Tape jobs typically request and mount tapes "by the ones". Each tape mount has a corresponding tape fetch. Modern robotic hardware minimizes fetch time when the volume is located inside the silo. When the volume is located outside the silo, slow human intervention is necessary. Most shops cannot afford to automate every volume and even the most astute use of silo content management techniques cannot ensure constant, cost-effective locality of reference. Manual tape handling is an unavoidable random access procedure.

We deployed the ThruPut Manager Robotics Setup Services (RSS) component to free the tape librarians. RSS ensures that all tape volumes are located where they are required prior to job initiation. Directed by console displays and/or picking lists, tape librarians fetch and enter tape volumes into the robotic silos. When all volumes are properly located, the job is allowed to initiate. RSS changes the manual, random access tape fetch into a sequential batch-like operation, significantly improving the efficiency of the tape librarian. The reduction of tape mount time to robotic speed minimizes job execution time. It also minimizes the amount of time that tape units are allocated and idle, thus maximizing physical and virtual tape hardware utilization.

### ***The Big Push...***

This was the turning point of our campaign. We were on the outskirts of batch supremacy. We were still outnumbered, but *they* were a largely ineffective and poorly led force. We had limited their ability to move jobs freely through the JES2 subsystem. We could run jobs, under our terms, in the place and time of our choosing. We had reduced them to singular guerrilla incursions, but they were still as dangerous as an unexploded mortar round. They had read-access to production data sets. They could still cripple us with allocation waits.

We began our big push on a cold Monday morning. We flew in low, under their radar. We had a secret weapon. We would hit them and they would never know they had been hit. It was a Psy/Ops operation. We deployed a package consisting of ThruPut Manager Dataset Contention Services (DCS) and

Ops/MVS automated operations scripts. It was a one, two and you're out, punch!

DCS controls and prioritizes access to data sets by batch jobs. It ensures that no job initiates unless it can allocate all required data sets. DCS has three rules of engagement: Standby, Contend and Claim. We use Standby service for non-production jobs and Contend service for production jobs. DCS is cognizant of a job's data set requirements from analysis of the job's JCL. It knows data set names and whether shared or exclusive control of the data set is required. Jobs with Standby service are allowed to run only when all data sets are available. Contend service is more aggressive. When a Contend service job requires exclusive control of a data set and that data set is in use, DCS dynamically queues for exclusive control of the data set on behalf of the job. When DCS receives control of the data set, it allows the job to initiate, passing control of the data set to the job during its allocation phase.

DCS has a sentry option to alert when a job initiation is delayed for data set contention. We enable this option for production jobs. When a production job initiation is delayed, DCS barks a multi-line alert message to the console. This message indicates which job is delayed, the data set names in contention and the holder(s) of the contended data sets. We have automation scripts that monitor this message. We scan the data set holder information and determine if the holder is production or non-production. We automatically cancel any non-production job that delays a production job. As part of the cancellation we post messages in the cancelled job's joblog indicating why the job was cancelled and the pertinent data set names. This product synergy is our missile defense system. Any problem allocation that is lobbed at us is identified and immediately shot down.

We also enabled a DCS option to order TSO users to free data sets when causing contention. The Nag option issues messages to TSO users signaling their culpability in the contention event, specifying the data set name and the name of the job requiring the data set. We extract information from the ACF2 LIDREC and augment the message with the name and phone number of the job submitter. This makes the contention hand-to-hand combat between users, taking us out of their hostilities.

### ***Debriefing***

We had nothing to fear, but fear itself. They gave us the world's best training and software and we made the best use of it. I have stories to tell my grandchildren. When they say, "what did you do in the war grandpa?" I won't have to say: "I sat in a stinking

cubicle and watched men die of thirst waiting all day for their jobs to finish."

We achieved all of our political objectives:

- We ended the anarchical control of initiators and tape units by aggressive users via judicious use of JLS limiting agents. All users now receive fair and equal access to computing resources. Control is fully automatic, yet flexible when circumstances dictate.
- We liberated the tape librarians and saved their valuable natural resources for their future by deploying RSS. Tape operators are more efficient, job elapsed time is minimized and tape hardware investment is maximized.
- Coalition forces secured and protected our production data sets from marauders through a combination of automation and DCS. Contention waits are reduced to near zero. Production data sets are made available to test jobs as much as possible so long as they do not interfere with production operation. Automation has proven itself a valuable ally in these joint operations.

Copyright © 2003 Kelly L. Vogt, All Rights Reserved. Permission is hereby granted to the Computer Measurement Group (CMG) to publish.

This work may not be reproduced by any means, in whole or in part, without the express written permission of the author or his lawful agent.