z - The Difference

BY DAVE RAFTEN

Many zSeries system programmers know intuitively that System z servers provide a more stable platform for business critical applications than other server platforms, and are constantly debating their co-workers. What is the technology behind making the "z" in "System z" stand for "zero down time?" This article will explain some of the features found in System z servers that specifically address availability.

What is availability?

The dictionary defines "availability" as the state of being ready for *immediate* use. This is how System z manages it. Yes, the server is unavailable during an unplanned outage, but it is also unavailable during planned outages such as for service or upgrades. That is why many of the RAS features of the server target the total picture, not just avoiding unplanned down time.

Concurrent maintenance

In the past, a common reason for planned outages was to apply hardware configuration changes. Now, virtually all of these requirements have been removed. In the past, whenever you wanted hardware maintenance done, you needed to power down the server and restart it with a "Power On Reset" (POR). Now, while production transactions are executing, you can replace and upgrade key internal components of the server. At the largest scale are "books." Each processor book contains CPs, memory, and I/O "fan-out" cards. On a minimum of a two book server, these books can be dynamically pulled and replaced for upgrades or repair while continuing to execute the workload. At a finer level, each individual I/O fan-out card within each book can be hot-pluggable without the loss of I/O connectivity. Every card in the I/O domain such as FICON port cards, OSA-Express, Crypto, and coupling links, and so forth. can be concurrently, repaired, or replaced. There is even concurrent power and thermal maintenance including the HMC and support element (SE). While this is happening, transactions continue without missing a beat.

Applying maintenance to software often requires the software to be "recycled," brought down and restarted, to pick up changes. On the other hand, System z servers are designed so microcode maintenance and even driver upgrades can be applied while applications continue to execute.

This is sometimes compared to changing an airplane's engine while the airplane is in flight, although a better comparison may be to not only change the engine, but also the wings, fuselage, and the cockpit.

Concurrent growth

Most capacity changes do not require an outage. You can grow the server from a single book sub-Uni 1-way to a 54-way dynamically to accommodate planned, or even unplanned capacity requirements. Emergency upgrades are possible with Capacity Back-Up (CBU) such as in a disaster recovery situation or loss of a server. Because many companies are subject to requirements to test disaster recovery capabilities, the CBU contract comes with five tests that can be renewed. CBU even supports specialized processors such as ICFs, IFLs, zIIPs, and zAAPs. For planned events, a planned upgrade can use the customer-initiated On-Off Capacity Upgrade on Demand (OOCUoD) offering. This allows one to upgrade to meet end of year capacity requirements, and downgrade the server in the first quarter! With the System z10 and z/OS R10, this is extended further with the Provisioning Manager. Rules can be set up defining when additional capacity should be provisioned to meet your business need. This provides a fast response to capacity and workload changes, and helps ensure processing power to meet your business requirements. More information on the Provisioning Manager can be found in the article “Don’t miss the goal! Capacity Provisioning can help” on page 63 of *z/OS Hot Topics Newsletter Issue 19,* August 2008, GA22-7501-15.

A balanced design requires memory and I/O to grow with the CP capacity. This is not a problem with the System z10 as one can concurrently add I/O and memory. After the I/O cards are added, you can then modify the I/O configuration definitions including channel paths, control units, and I/O devices. You can also add or remove LPARs to a new or existing logical channel subsystem, and then dynamically add cryptographic features to existing LPARs. While you are changing the airplane’s engine, wings, fuselage, and the cockpit while in flight, you are also growing it from a simple two-passenger plane to a 400 passenger luxury-liner!

A fault tolerant design

A *fault tolerant* design allows the system to continue running if there is the loss of a single component. Moreover, it is designed to do this without even impacting transactions. System z provides fault tolerance for all of its key components. This includes not just the CPs (transparent CP sparing), memory (dynamic memory sparing), or I/O (I/O Interconnect), but also the timing oscillator card, power supply, channel paths, OSA cards, support elements, and others. Through internal monitoring, possible problems are detected and problem components are designed to be switched over without even failing a single transaction.

This is in addition to IBM’s most robust processor design with intelligent retries and a very elaborate and successful design of internal recovery hardware.

Server phone home

In a certain movie, an extraterrestrial was able to “phone home” to communicate to his support crew for help after he needed it. Using somewhat better technology, IBM computers are able to detect potential error situations *before they become problems* and phone home (including web-enabled communications), calling IBM so the CE can schedule time to come and replace the potential problem part. This is fully automatic, done by the hardware, while service can be done without impacting application availability such as z/OS or Linux. This long standing design feature starts dispatching parts and personnel immediately after a machine event occurs.

Application availability

The benefits discussed are available on any application that can run on the System z10 platform from day one. This includes not just z/OS, but also z/TPF, z/VM, z/VSE, and Linux on System z operating systems. The hardware does this without operating system interaction.

If you are running z/OS, there are additional RAS features that do require hardware interaction. One example is the *CICS subspace group facility*, which provides application availability to protect the storage allocated by a transaction from being read or overwritten by another transaction, even from within the same address space. As an additional benefit, this also provides increased security, protecting confidential data from being maliciously read by an unexpected program.

Application availability also means protecting applications from malicious attack. To support this, System z processors support many industry standard cryptographic algorithms by using imbedded hardware for additional security as well as performance. Support includes AES up to 256 bits, SSL acceleration for z/OS and Linux, Triple DES, SHA-256, and others. System z has been certified with Crypto Express2 being FIPS 140-2 Level 4 certified. LPARs are Common Criteria Evaluation Assurance Level 5; the highest possible rating.

While there are many more standard availability features, this can explain how System z servers deserve their “z”, designed for Zero down-time.