HIGH AVAILABILITY FOR THE MAINFRAME ENVIRONMENT

GLOBAL TECHNOLOGY SOLUTIONS GROUP

MAINFRAME MANAGED SERVICES

In this paper, we provide an overview of High Availability for the mainframe environment.

To help set the context, we're going to use the perspective of a Financial Services institution running on z Systems. The firm's business customers and relevant regulatory agencies impose stringent requirements for continuity.

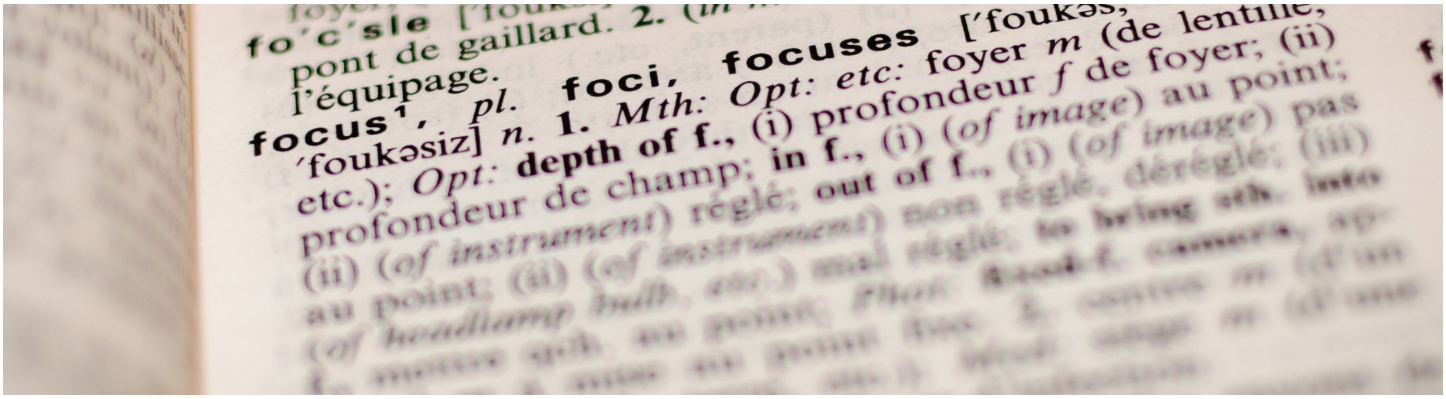

TERMINOLOGY

HIGH AVAILABILITY (HA)

- The elimination of single points of failure within the environment
- The ability to recover a failed unit of work (web connection, transaction, batch job, DDF connection, etc.) on the surviving member(s) of the HA environment
- The ability to sustain outages in a member of the HA environment without impacting delivery of service to the enterprise
- The ability to provide continuity of service to the enterprise in all situations short of a disaster recovery declaration

Horizontal Growth: adding capacity to the environment in the form of new processing footprints, i.e. adding new CECs to the environment

Vertical Growth: adding capacity to the environment by increasing the computing capacity of the existing platform footprint, i.e. adding new engines to existing CECs

Sysplex HA components: Those components that, when implemented, provide a direct HA function and are required building blocks of the HA environment

Sysplex Enablers: Those components that, when implemented, support Horizontal Growth by enabling sharing of network, data, etc.

HIGH AVAILABILITY IS A CAPABILITY

HA requires a review of the business applications considered critical and a determination of which might require infrastructure alternatives or application modification.

High Availability is not a "product." It is a capability achieved by a holistic approach to the implementation of numerous technologies, the adoption of future state application development architectures, and the implementation of a plan for horizontal growth.

HA requires a review of the business applications considered critical and a determination of which might require infrastructure alternatives or application modification. Maintaining a state of high availability also requires the adoption of regularly practiced exercises, including failover, giveback, takeback, diligent documentation, and most importantly, recovery exercises.

HIGH AVAILABILITY CONTEXT ON Z SYSTEMS

WHAT WE CAN PROTECT

The Point of Entry (POE) into z Systems is generally "protected" by HA techniques. This includes IP connectivity handled by Sysplex Distributor. Sysplex Distributor maintains and presents outward visibility of z Systems IP space – meaning that the HA environment continues to present consistent IP visibility to the outside world.

Inside the POE, the implementation of HA techniques creates an environment capable of surviving

- CEC failures (assumes a pair of CECs in the HA footprint)
- DASD subsystem failures (this requires Metro Mirror with basic Hyperswap)
- LPAR failures
- z/OS failures
- Individual ASID failures (CICS, DB2, MQ, IMS to cite a few)
- Scheduled outages of the items listed above

WHAT WE CAN'T PROTECT

z Systems HA does not provide any improvements to resiliency outbound from the Open Systems Adapter (OSA). In other words, if the next hop out of the mainframe network goes to a switch or router that has a single point of failure, a failure at that point will potentially render the z Systems environment inaccessible because of the failing switch or router.

Just like a network exposed to single points of failure, z Systems HA techniques can't protect a storage topology that is not resilient. However, some HA techniques are available for use (Basic HyperSwap) which will improve the resiliency of the storage subsystem when it is properly provisioned.

We always recommend that the network, SAN fabric, FICON directors and such be provisioned so as to mitigate any single point of failure.

Network and Storage resiliency, outside the z Systems boundaries, is a separate topic that we will address in future white papers.

HIGH AVAILABILITY IS BUILT ON THE FOUNDATION OF PARALLEL SYSPLEX COMPONENTS AND ENABLERS

HA is made up of components and is facilitated by enablers.

We can think of these required **components** as the "foundation" for the HA house:

- Sysplex Distributor
- VIPA/DVIPA implementations
- Coupling Facility used by all HA components
- Coupling Facility 'self-healing' in terms of duplexing, sizing, reconfiguration
- VTAM Generic Resources and Multi-node Persistent Session

The **enablers** below are some of the techniques that leverage the foundational items above. Each of the items below offers a product-based solution for high availability.

Each of the items in the list below requires that some or all of the foundation components be in place. In the case of the list below, all are classified as Data Sharing technologies.

- DB2
- IMS
- MQ
- CICS Temporary Storage
- VSAM RLS
- Transactional VSAM
- Others

ESTABLISHING THE HA ROADMAP

We first define what we're trying to protect: is it a single workload; multiple workloads; or pretty much everything? It's best to involve the business units – to gain their insights as to their current and future state requirements.

Then, we set some of the parameters.

- What is the tolerance for scheduled outages? More and more we're hearing "no outage – regardless" from the business.
- Perform an outage analysis – for both planned and unplanned events – and develop mitigation actions for each outage occurrence.
- Is there a mapping of the applications to infrastructure components that provide highly available configurations?
- Is the scope a single site, or multiple sites?
- What high availability "enablers" are we going to implement? In what order?
- What outages did you discover that can't be fixed with traditional high availability?

We need to determine the optimal sequencing. Most items will require serial implementation; the balance should be sequenced according to goals, requirements, budget, resource availability for testing, etc.

As with any significant initiative, we need input from both the technology and the business units. It's a mistake for a technology group to try this in a vacuum. Rather, we must be sure to involve those within the Application Development groups that are responsible for defining the future-state application development patterns.

Finally, we aggregate high-level costs with the HA Roadmap for budget and staffing approvals.

BUILD THE REQUIRED FOUNDATION

We advise our clients to use a fully outfitted test environment. This will be your HA staging area for the future. Plan to devote resources to it. The key to successful testing is to have a representative test environment available.

In our experience, it's easier to start at the "outer edges" of the z Systems and work inwards:

- Get the IP connectivity working for failover
- Get the workload distribution working for horizontal components
- Test failover and takeback
- Test until confident
- Document, document, document
- Automate where required and where possible – but integrate into your documentation what automation is really doing

We advise the use of TN3270 server(s) as playground for testing and recovery. We also validate that Workload Management (WLM) is involved in workload distribution decisions.

BUILD THE ENABLERS

Next is to implement the "enabler" components – DB2 Data Sharing, MQ / CICS / RLS – the data sharing techniques – and Port Sharing (same LPAR) which is frequently used for CICS. Then, we document possible recovery scenarios, simulating the failure of individual component – LPAR, CF, CICS, DB2 data sharing group member as examples.

We are then ready to define the precise commands utilized for recovery, status assessment, etc., defining automation to speed the recovery of transient situations. Finally, we define the application landscape (DB2, DDF, CICSPlex, application affinity considerations, and other considerations) and how that is supported by our availability solution.

OPERATIONAL CONSIDERATIONS

Visibility: The current state environment (today) makes it easy to tell "what runs where." The future state environment will be flattened horizontally and at any one point in time, it might be less than obvious where a given unit of work is running.

Education and training: Operations personnel will need to understand future state topology, and will need to be mentored on new messages, commands, and other operational aspects required to monitor and manage the environment.

Operations will have a role in the handling of planned outages on a "member by member" (of the Sysplex) basis. This will be a new activity for them to perform. Automation will most likely mitigate the complexity. Nonetheless, quiescing a single member and then introducing it back into its original role after the incident has been fully diagnosed and service is restored is not a simple task.

Our advice is for Operations to work together with the Technical and Application teams in the testing environments. This gives both of these teams a view of how all the moving parts work together, providing valuable experience.

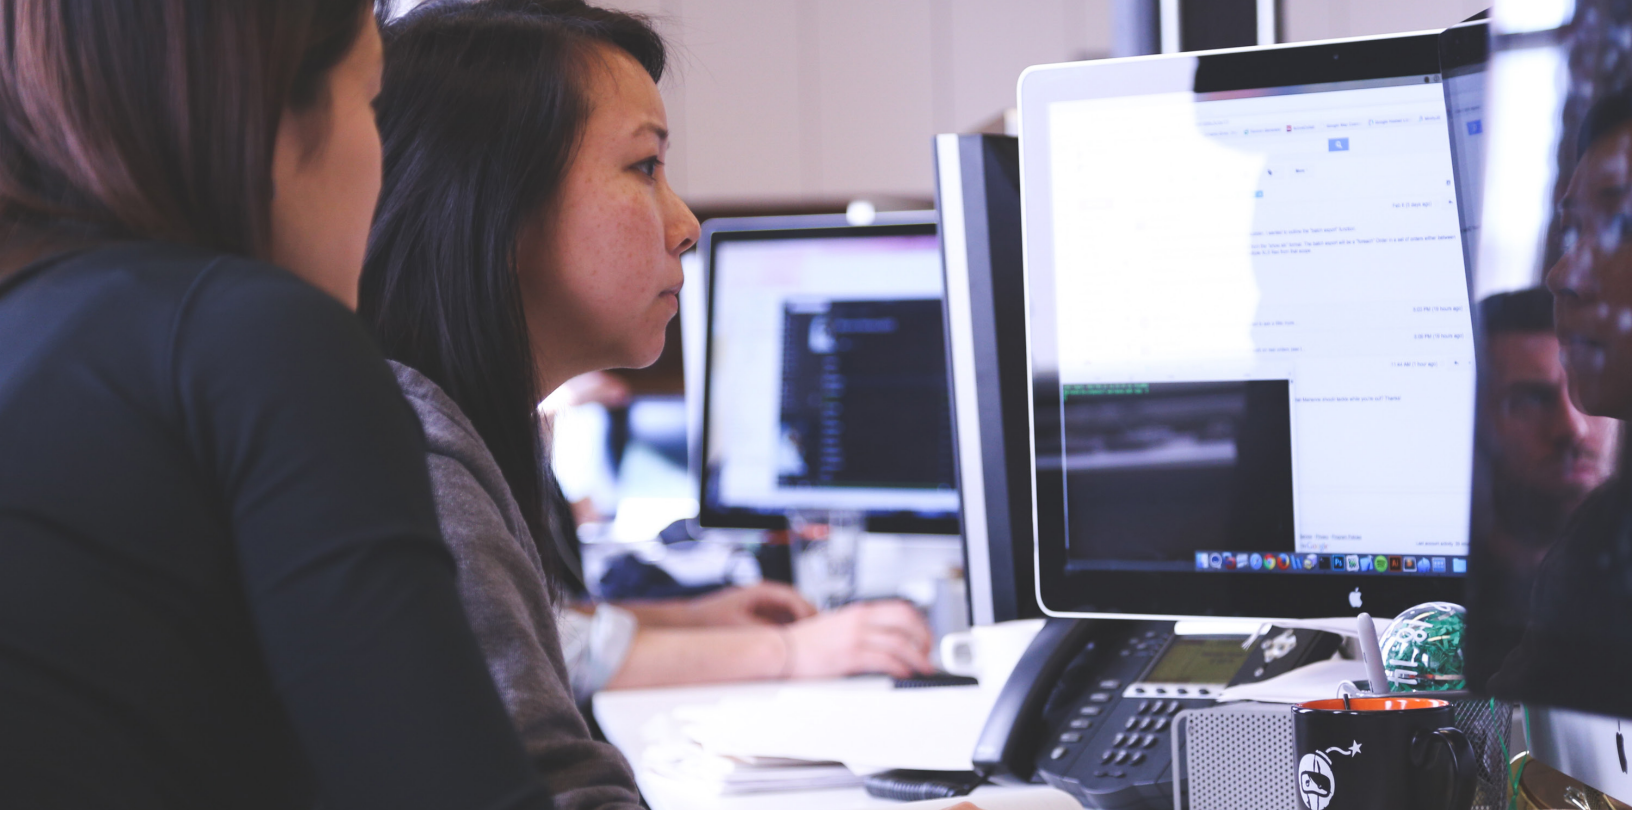

HIGH AVAILABILITY REQUIRES EDUCATION

HA will introduce a new level of complexity that can only be mitigated, to a degree, by training your staff.

The infrastructure you build will become a way of life – everyone needs to be on the same page. HA will introduce a new level of complexity that can only be mitigated, to a degree, by training your staff. Depending on the preferred mode of training at your firm, you may choose among formal classes, informal training sessions, or a combination thereof.

"Formal" classes cover the full breadth of the particular technology. The downside is there is a lack of opportunity to customize. Budget issues **may prevent the entire staff from participating in the training**, and it operates on a set schedule which may or may not be convenient.

"Informal" sessions are focused training for your environment. They are customized, and provided "just in time." They leverage the test environment for exercises. If delivered via WebEx or the like, they can be recorded for future use.

You can also synthesize the best of both worlds. If a formal class is utilized, have the participants provide training to others. Some technology will not warrant a formal class, so deliver this training via informal sessions.

Again, training should be timed to coincide with adoption of the technology… so-called "just in time" training. Participants should include Technical Support, Operations and Application Architects.

HIGH AVAILABILITY LEVERAGES WORKLOAD MANAGER DECISION MAKING

Addition of an additional CEC affords a way to become tolerant of HW-based outages (planned or unplanned).

WLM plays a role in workload distribution at Point of Entry. Sysplex Distributor utilizes WLM decisions in determining which LPAR to route a given request to.

Leverage benefits of Intelligent Resource Director (IRD) for LPAR-level balancing. Think of IRD as "WLM for the CEC."

Scale by adding capacity (CECs, LPARs, regions/subsystems) horizontally. Growing the Sysplex with an additional CEC affords a way to become tolerant of hardware-based outages – planned or unplanned.

Some components, like TN3270 and others can be **quiesced to "shift" workloads** to other LPARs or CECs.

Takeovers can be invoked manually during low-volume period to lessen impact.

HIGH AVAILABILITY APPLICATIONS

Applications can run in an HA environment as they do today – with a gradual migration to a "run anywhere" mode.

Everyone always worries about the applications. We typically hear:

- *"My application isn't Parallel Sysplex enabled"*
- *"My application can't run in two places at once"*
- *"My application has to run in a particular CICS region"*

Take a "fact or fiction" approach:

- Not all applications will **leverage** Parallel Sysplex capabilities
- Nearly all applications will **benefit** from Parallel Sysplex capabilities

Very old legacy transactions may have hard-coded dependencies to a given CICS region – in our language, these transactions have an affinity to a given CICS. These include:

- CICS APPLID name coded in programs
- Home-grown data in memory applications that store counters and such
- Other techniques of storing session state data that cause an affinity

Applications can run in an HA environment as they do today – with a gradual migration to a "run anywhere" mode.

We recommend leveraging CICS Parallel Sysplex techniques – temporary storage server, named counter, Data Tables, etc.

Affinities can be addressed within CICSPlex. They can be remediated on a case by case basis if the ability to "run anywhere" is deemed to be worth the remediation effort.

It is possible that some applications will not "play well" in a DB2 Data Sharing environment. These include transactions or jobs that are long running processes without frequent commit logic, or those that introduce more locking than required for the specific use. Issues here will manifest as performance problems; a good DBA review is required.

HIGH AVAILABILITY MAY PRESENT A CHALLENGE

Individual components, implemented without the vision, tend to be exercises in wasted time or set poor perceptions.

Don't try to turn HA into a DR solution. HA is about redundancy and fault tolerance – it does not improve your DR capabilities. In fact, it can complicate them.

It's not a quick fix – be prepared to manage it as a strategic initiative. Deploy according to a plan – **not as a series of knee jerk reactions to outages.** You need a holistic view and plan. Individual components, implemented without the vision, **tend to waste time or set poor perceptions.**

Education. Make that a component of the HA Roadmap so the team realizes you recognize the need. Provide it – but in a just in time manner. The breadth of technology is wide indeed, and drinking from the education "fire hose" will not be effective.

Applications. Not all applications play well in data sharing modes. Not all applications can be split across multiple regions without affinities. Don't let these facts derail the HA plans – there are ways to handle these challenges.

HIGH AVAILABILITY SUCCESS DEMANDS A PLAN

Don't go out and get creative with some complicated recovery scenario unless it's to protect a mission critical application and no other component can provide the same capability.

The first step is to set goals: determine what you are trying to protect with HA.

When we design, we assess the long-term role of additional CECs and the role of secondary sites and how HA implementation might impact those environments. We then develop the HA roadmap, build the required foundation, and build the "enablers" into the environment so they are staged for production deployment.

We recommend that you build the capability but do not utilize it – yet. Build everything you need and then start a regimented testing plan - go slow. We then proceed to:

- Identify your pilot test cases and your "what if" failure scenarios.
- Define your infrastructure testing criteria.
- Define your performance baseline & benchmark cases.
- Define your documentation and automation guidelines.

From there, we develop the implementation plan.

Ban "one off" implementations: stick to a service catalog approach where standard high availability offerings are defined and mapped to use case(s). Using a standardized approach keeps the high availability scenario as simple as possible.

If a high availability solution is required that isn't provided by an IBM or ISV offering, then integrate it into the Service Catalog as an acceptable recovery solution – but clearly document the isolated use case(s) to which it applies.

Implemented properly, High Availability can be the game-changing process which meets the resiliency demands of today's "always-on" world.

GTSG has been supporting the mainframe since the 1980s. We'd be delighted to discuss further: please write us at mainframe@gtsg.com or call 877-467-9885.

GLOBAL TECHNOLOGY SOLUTIONS GROUP

T 877 467 9885
F 877 225 4084
W gtsg.com
E mainframe@gtsg.com