Parallel Sysplex: Case Studies

OS/390 and Storage Systems Technical Conference
Como, May 2000

Paul Arnerich
Tsd (UK) Ltd/Laird Ståhl Ltd
Paul_TSD@compuserve.com or Paul@tsdd.co.uk

Agenda

- Sysplex Review
  - → Basic
  - → Parallel
  - → Application Issues
- Business Issues
  - → Benefits
  - → Justification
  - → PSLC
- Case Study 1
- Case Study 2
- Where To From Here

Sysplex Review

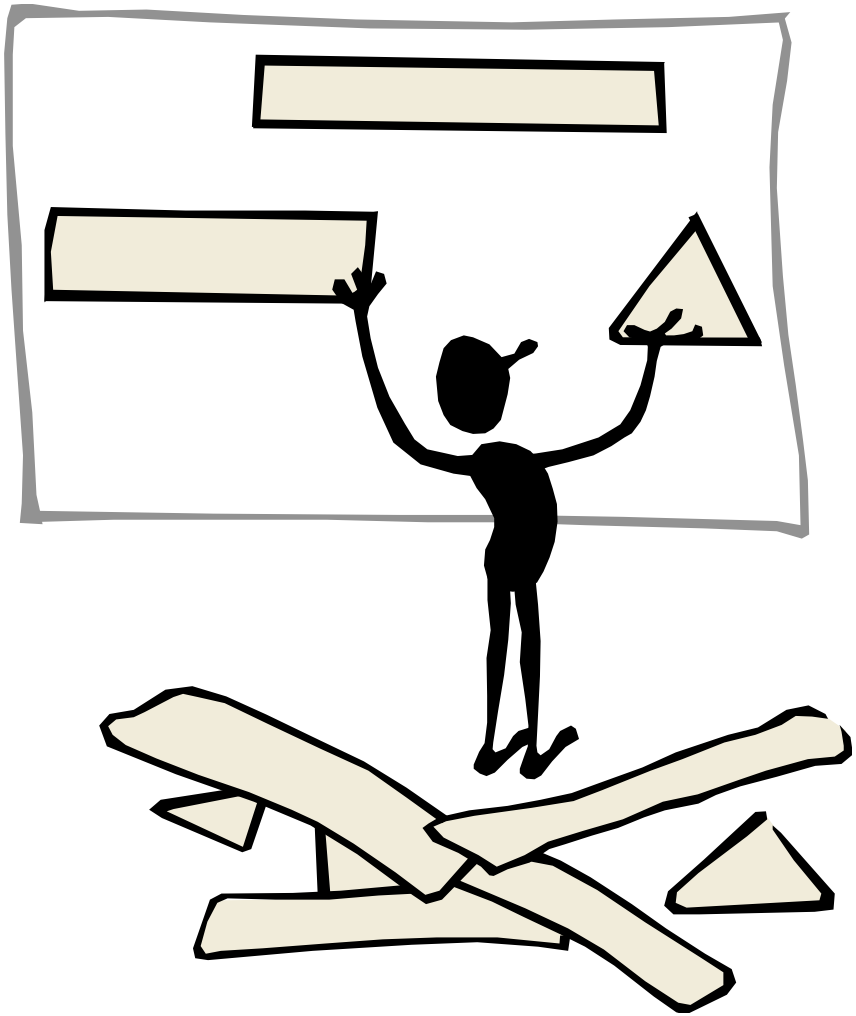

Basic Sysplex

- **Communications**
  - → **XCF - OS/390 base software to provide high speed communication protocol between MVS's**
  - → **Simple API available to any authorised assembler program**
  - → **Provides Messaging services**
  - → **Provides Status Monitoring services**
  - → **Needs a physical transport .... a CTC**
- **Sharing**
  - → **GRS must propagate system locks to participating MVS's**
  - → **GRS can propagate dataset enqueues to participating MVS's**
  - → **Can share DASD resources with integrity**
  - → **Must share selected systems datasets**

Basic Sysplex ...cont

- **Time Base**
  - → **"Time is an illusion, lunchtime doubly so"**
    - **- Ford Prefect (Hitchhikers Guide to the Galaxy)**
  - → **Must have a common time reference**
    - **used for logging for recovery purposes**
    - **used as a measure of health**
- **Limitations**
  - → **No more than 32 MVS's in a single configuration**
  - → **GRS propagation can cause significant performance overhead**
  - → **Tuning of GRS required if more than 4 MVS's**
  - → **Consider GRS STAR configuration if over 4 MVS's**

Basic Sysplex

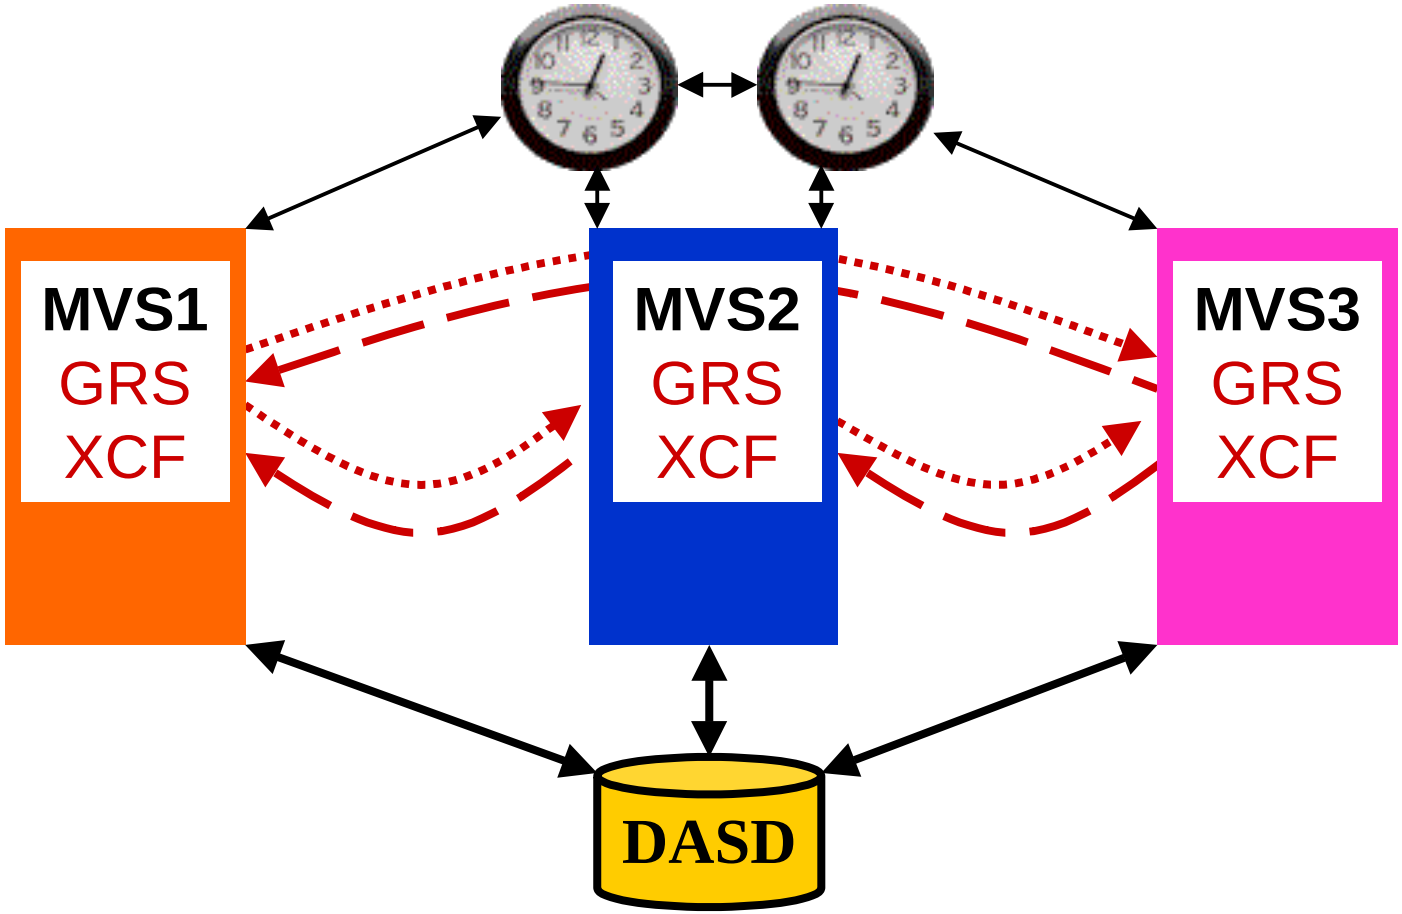

Basic Sysplex - Pre-reqs

- 2 or more processors
- 2 or more shared DASD volumes
  - → times two
- Sysplex Timers
  - → times two
- CTC's providing any to any connectivity
  - → times two
- Appropriate OS/390 Software levels
- Applications to exploit the architecture

Parallel Sysplex

- Adds to Basic Sysplex through provision of shared memory
- S3 - Shared Standalone Storage
- Allows true data sharing between MVS's as if it was processor memory
- Delivered via the Coupling Facility:
  - → A Stand Alone Coupling Facility, or
  - → Internal Coupling Facility
    - an engine on a CMOS processor
- XES - Base OS/390 software
  - → provides all communication with Coupling Facility
  - → Simple API available to authorised assembler program

Parallel Sysplex

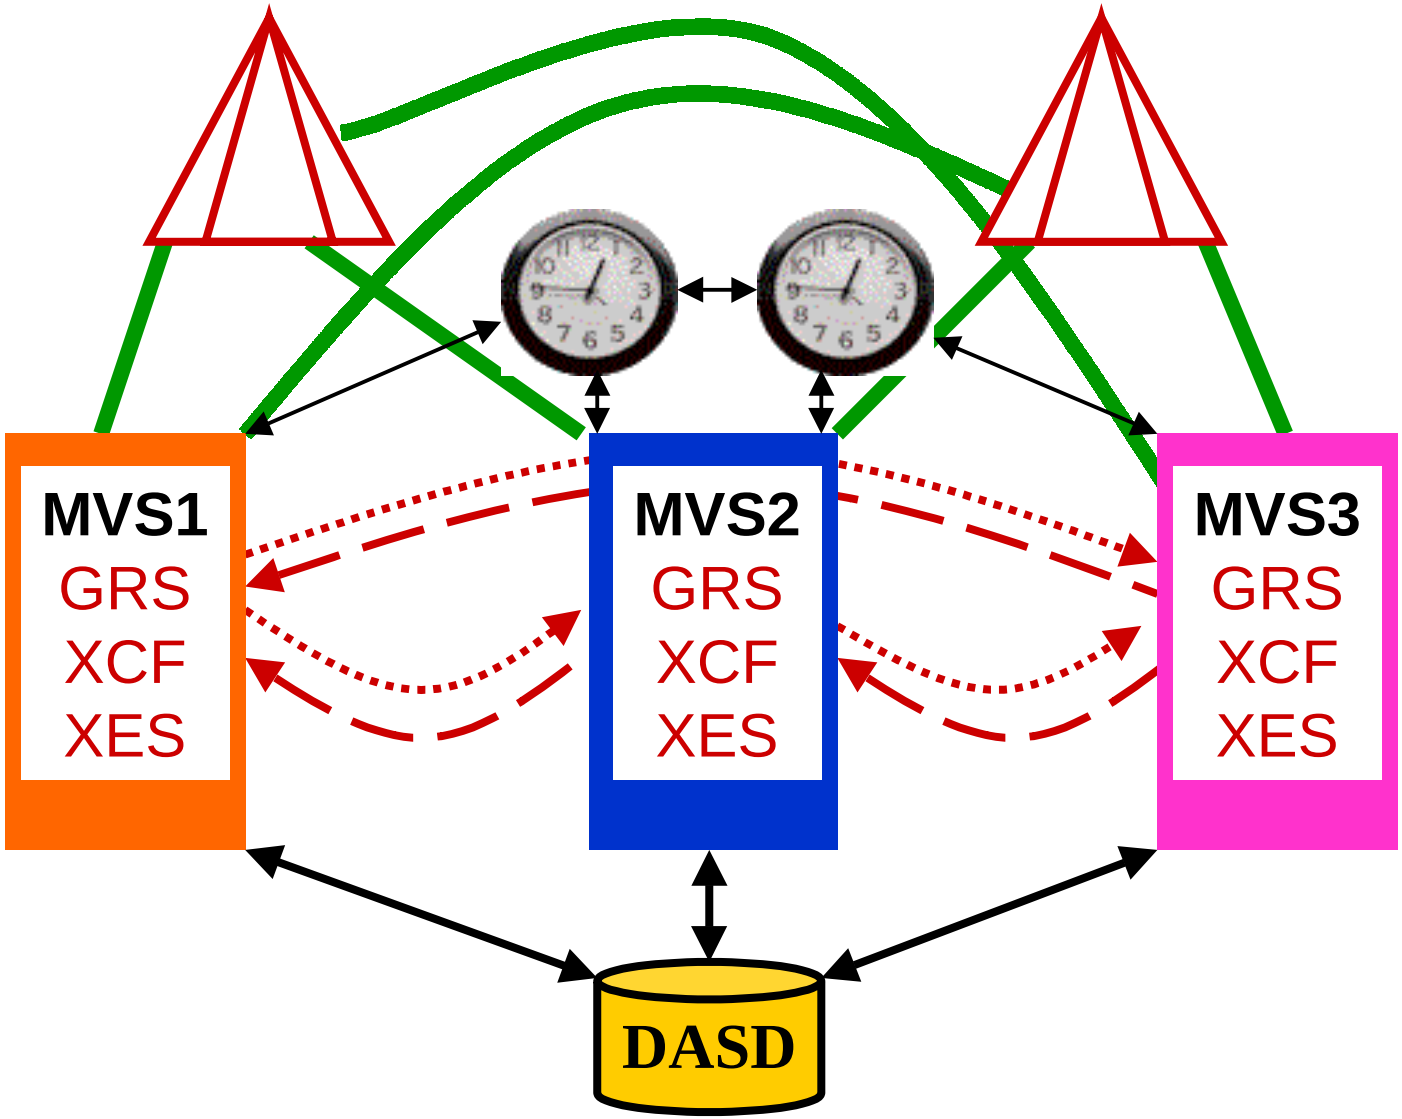

Parallel Sysplex - Pre-reqs

- **Basic Sysplex**
- **Coupling Facility**
  - → **times two**
  - → **sufficient storage and Mips to do the job**
  - → **Scalable**
- **Coupling Links**
  - → **times two**
  - → **Currently 100 Mps, 250 Mps or 1000 Mps**
  - → **need sufficient bandwidth**
- **Appropriate OS/390 Software levels**
- **Applications to exploit the architecture**

Supported Applications

- XCF
  - → MCS
  - → GRS
  - → RTM
  - → APPC
  - → JES2
  - → JES3
  - → DAE
  - → VLF
  - → SMS (PDSEs)
  - → TSO/E
  - → OPC (Tivoli)
  - → WLM
  - → CICS & CTS
  - → etc
- XES
  - → RACF
  - → JES2
  - → JES3
  - → GRS
  - → DB2
  - → IMS
  - → DFP (Tapes)
  - → CTS
  - → VSAM
  - → VTAM
  - → etc

Warning

- Application design is critical
- Choice of Application Enablers is critical
- Which CF holds which structures ?
- Co-location of CF with Application Enabler is essential to obtain best access speeds

- Ensure Coupling Link bandwidth is sufficient
- H/W design requires duplicity of critical components

- Investment in operational infrastructure is essential

Business Issues

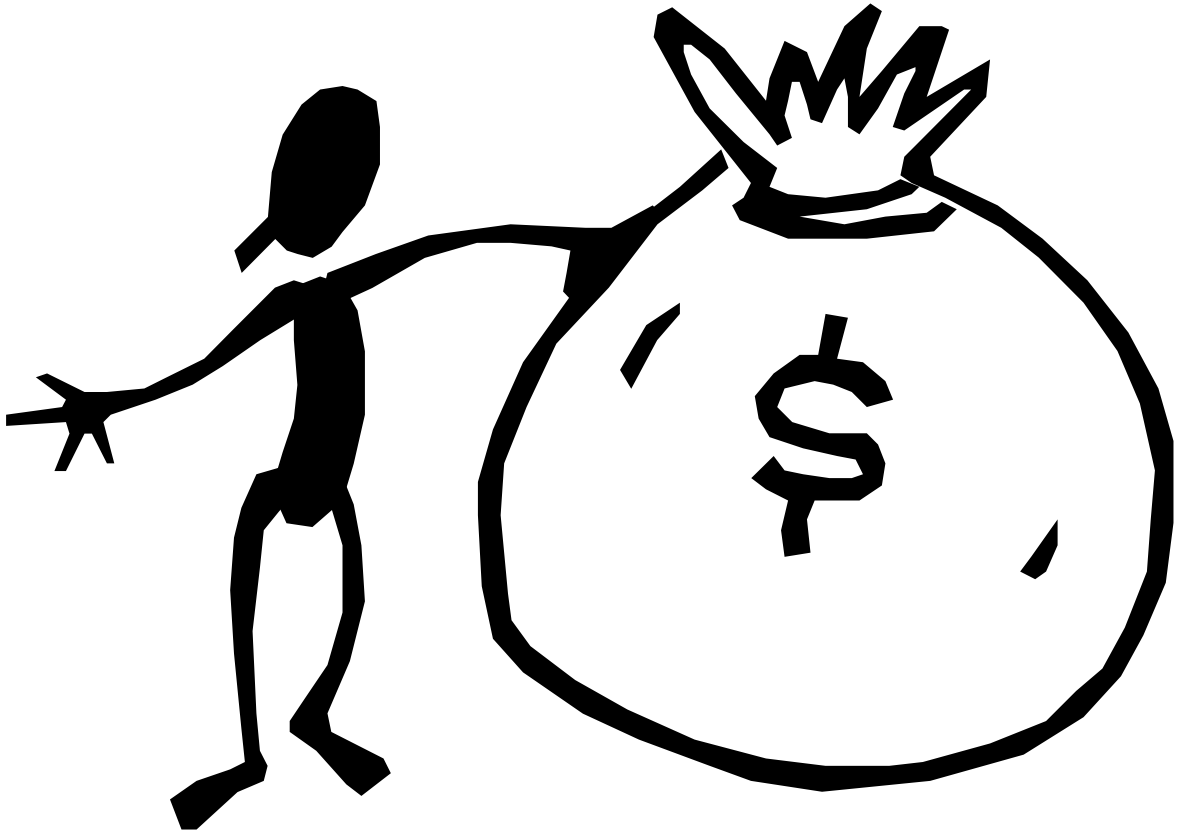

Business - Benefits

Parallel Sysplex benefits when fully implemented

→ *Rapid response to unexpected growth*

- **Quickly add power to match requirements without disrupting business**

→ *S/390 Resource Sharing*

- **Value today in a single CEC and multiple CEC environments**

→ *Balance multiple workloads*

- **Let all applications share system resources in order to meet business goals *you define***

→ *Increased productivity*

- **Manage multiple systems as single system from single point of control**

→ *Continuous application availability*

- **Leading-edge application availability**

→ *Investment protection*

- **Build on current investments in hardware, software, applications and skills, all at a reduced cost of computing**

Business Justification (1)

- Need a reason to implement
- Requires acquisition of hardware
- Requires investment in operational infrastructure
- Requires investment in application design

- So why do it ?

Business Justification (2)

- **Two compelling reasons for Parallel Sysplex**

1) Save money (Expediency)

→ **Parallel Sysplex Licence Charging (PSLC)**

→ **SHAMPlex**

→ **S/W TCO reduction**

- **Probably needs aggressive Asset Management**

2) Availability

→ **99.999% availability is possible**

→ **Evaluate availability requirements**

→ **Does loss of service means loss of revenue ?**

PSLC requirements

- **PSLC allows customers to qualify for significant reduction in OS/390 licence charges**

- **To qualify the customer must have:**
  - → **Functional Parallel Sysplex**
  - → **50% of 'production workload in said Plex**
  - → **50% of Mippage in said Plex**
  - → **Plex must be exploiting the Coupling Facility**
  - → **Valid exploiters range from full DB2 data sharing to the sharing of a single Tape Drive**

Availability Definitions

- **High Availability**
  - → **Provide access to applications during planned outages**
  - → **Achieved through use of redundant components and thorough testing**
- **Continuous Operation**
  - → **Provide 'service' at all times - no outages**
  - → **No change activity is possible**
- **Continuous Availability**
  - → **Blend of the two**
  - → **Continuous access to applications, no outages**

Case Study 1

Case Study 1 - Background

- Insurance and Financial sector
- Historical configuration of two processors due to acquisition
- Not a shared or integrated workload
- Needed Mippage for Y2K development
- Lip service only to high availability
- Software costs growing rapidly
- Feasibility study to verify costs of Parallel Sysplex implementation versus PSLC savings
- SHAMPlex an obvious choice

Case Study 1 - Business Case

MVS S/W costs p.a (IBM only)	**600 groats ***
H/W (CF, Clinks, CTC's, Dasd)	**200**
Internal Resources	**35**
External Resources	**7**
P.U.	**8**
Year 1 PSLC Savings	**50**
Future Year PSLC Savings	**250**

* NB: not real currency values, but in proportion

Case Study 1 - Implementation

- Acquired small CMOS processor, two second hand Timers, CTC's
- Already at appropriate OS/390 level
- Parallel Sysplex covered:
  - → 2 Production images
  - → 1 Coupling Facility
  - → 2 Development images
  - → 3 Y2K development images (not time sensitive)
- Included Y2K and Develpoment MVS's to qualify on 50% rules
- Sharing most Cartridge Drives
- Minimal operational changes by design

Case Study 1 - Summary

- **No Sysplex side benefits:**
  - → **Console consolidation**
  - → **Shared Spool**
  - → **Data sharing**
  - → **ISV Software savings**
  - → **etc**
- **Just overheads:**
  - → **Increased operational complexity**
  - → **Slight performance overhead of XES, XCF, GRS**
  - → **Technical Staff frustration**

Case Study 2

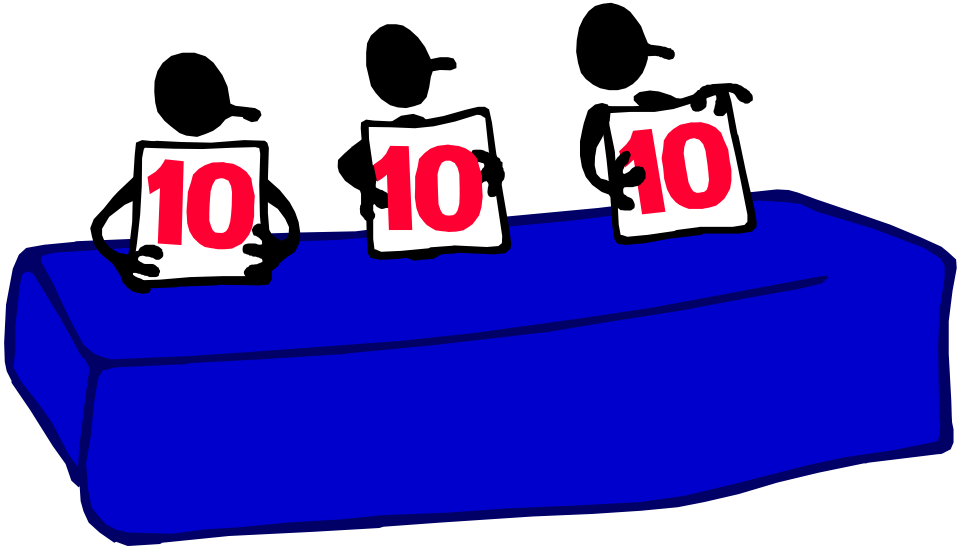

Case Study 2 - Background

- Large European Airline
- Large Mippage required for DB2 based reservation system
- Size due to growth
- Aging (creaking) Bipolar processors, not owed anything
- Replacement CMOS processors could satisfy most of the Mippage
- Availability crucial from business perspective
- Application design modern
- Full DB2 data sharing Parallel Sysplex an obvious choice

Case Study 2 - Business Case

- **Simple !**

- **Loss of reservation system meant loss of income immediately**
- **Long term loss of reservation system meant loss of organisations viability after ‘n’ hours !**
- **Processors due to be replaced imminently**

Case Study 2 - Implementation (1)

- **Acquired three CMOS processor**
  - → **Sufficient Storage for large DB2 Caches**
  - → **Engines specifically for Internal CF usage**
  - → **Highly available hardware**
- **Acquired High Band width CLinks, CTC's and Timers**
- **Already at appropriate OS/390 level**
- **Parallel Sysplex covered:**
  - → **3 Production images**
  - → **3 Coupling Facilities**

Case Study 2 - Implementation (2)

- Single system image of the processing capacity
  - → JES MAS
  - → Sysplex operational console approach
  - → Shared RACF
  - → Fully shared DASD
  - → etc
- Migrated existing DB2 Applications to new platform
- Other applications followed
- Location of CF's to DB2's is critical from performance perspective
- Care taken to minimise Inter System Read/Write Interest

Case Study 2 - Summary

- **Implementation costs balanced against delivery of highly available system**
- **Significant Operational benefits**
- **Performance costs insignificant - due to design**
- **Will acquire higher bandwidth CLinks as they become available**
- **Achieving significant reduction in ISV software costs by licensing only for capacity required**

Where To From Here ?

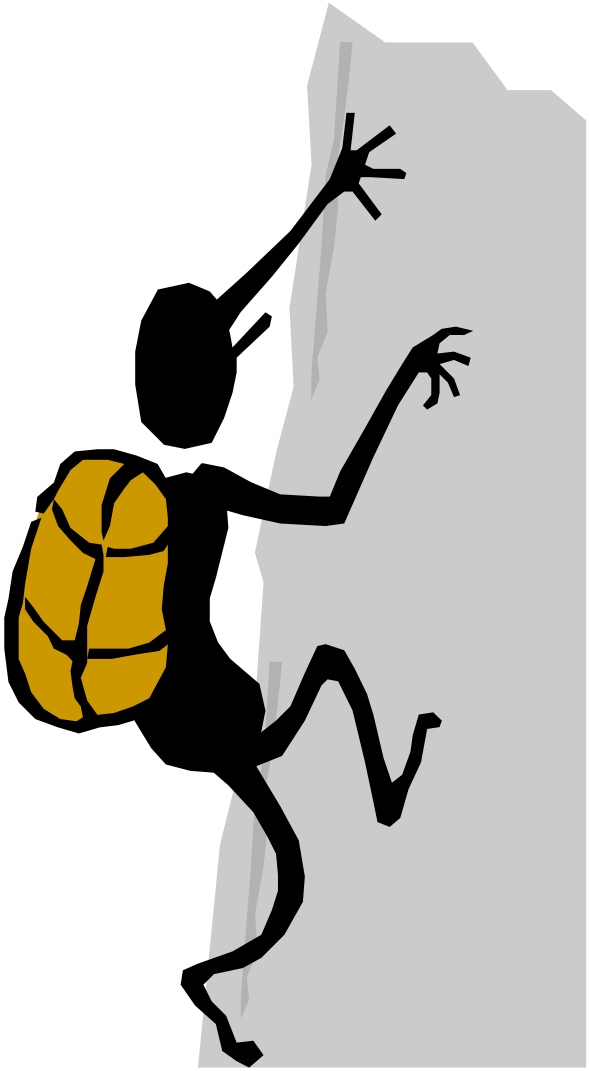

Operational Infrastructure (1)

- **Capacity management**
  - → **Monitor CF and CLink usage**
  - → **Understand growth requirements and consumptive spikes.**
- **Change management.**
  - → **Some Sysplex changes are introduced across all images**
  - → **Understand when this is required and avoid where possible**
  - → **Dynamic Reconfiguration should be used in every case.**
- **Problem management.**
  - → **Some problems on one MVS can affect the viability of all MVS's, from a Sysplex health perspective**
  - → **Ensure CF dumps are trapped, and saved**

Operational Infrastructure (2)

- **Availability management**
  - → Ensure configuration allows for n+1 redundancy in all component parts
  - → Ensure (through trial/drill) that the important MVS's can sustain loss of Sysplex components
- **Audit and controls**
  - → Naming conventions used for 'things' in the Sysplex must be documented and understood
- **Operations Management**
  - → Many Sysplex failures are self correcting
  - → They do produce 'frightening' messages
  - → Operational Drills are essential

Summary

- **Continue with TCO reductions**
- **Geographically Dispersed Parallel Sysplex**
  - → **Use PPRC to mirror DASD across sites**
  - → **Cross Locate CF's**
  - → **Watch distances, performance implications**
- **Requires highly available applications**
- **Requires investment**
- **Practice failure**
- **99.999 is a reality if desirable**
- **Is 99.999 really desirable ?**

Questions

???