

Intelligent Monitoring System (1987-Present)

THEN

AND

NOW

W. E. Farrell
SAIC

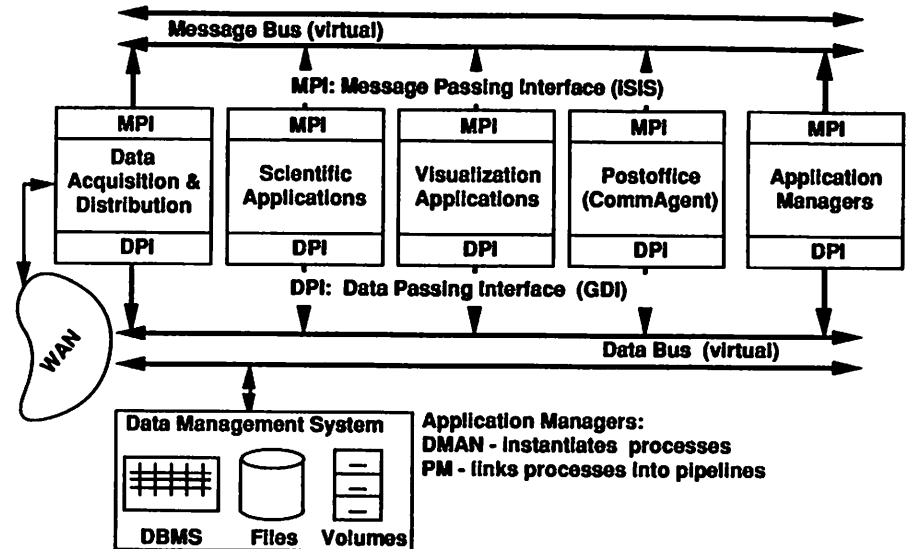
NOW Workshop
ASPLOS-VI
October 4, 1994

- Seismic data processing:
 - Time-series in (1-10 GB/day)
 - Earthquake list out (1 MB/day)
 - Support ad hoc research
- World-wide data collection
 - ~50 sites
- Analysis at Center for Seismic Studies,
Arlington VA
- Cornerstone of the International Data Center
of the U. N. Committee on Disarmament.

Overview

- ~ 20 workstations on an Ethernet
- ~ 50 GB magnetic disk
- ~ 300 GB Epoch optical jukebox
- ~ 600 kloc
- ~ 40 loosely-coupled Unix processes
Some stand-alone, many interactive
High degree of parallelism
- Unix, Isis, Lisp, (Oracle, Ingres, Sybase)

Architecture



Data Passing Interface

- Our original API required two functions (get and put) for each table of the schema. These contained a small amount of embedded SQL


```
get_Table1(tableName, whereClause, rowPointer, ..)
put_Table1(tableName, whereClause, rowPointer, ..)
```

- The Generic Database Interface is more powerful

```
gdi_submit(database, SQL, rowPointer, ..)
```

- Supports joins and abstract data types
- Works with 4GLs (IDL) too

Message Passing Interface



CommAgent — Communications Agent

- Provides logical name-space service
 - Maintains message queues
- DMAN — Distributed Application Manager
- Creates processes to satisfy service requests

Message Passing Interface

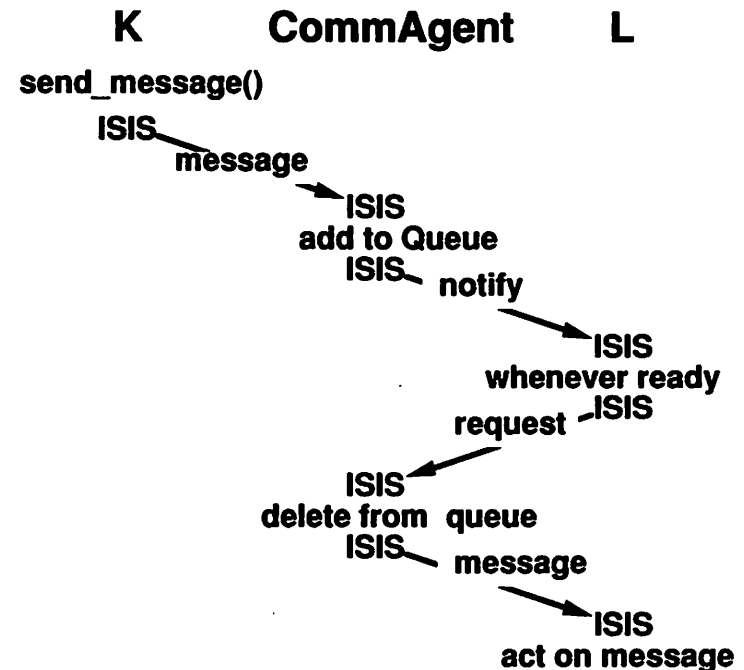
(continued)

ISIS

- Monitors processes
(~ few second response time)
- Process groups
(CommAgent == group of 1, but with many clients)
- Shadow CommAgent and DMAN for fault tolerance (future)

Message Passing Interface

(concluded)



- If a queue for process L does not exist, the notify alert goes to the DMAN application

Job Control

Interactive Applications

- User's serial activities request services from remote processes, as appropriate.
- The IPC triad (ISIS/CommAgent/DMAN) ensures server processes are always available.

Batch Applications

- Numerous applications do not have built-in IPC.
- For these a "wrapper" program that speaks IPC is written to command them
- Maps IPC messages to command line args.

Job Control

(continued)

Process Pipelines

- Process Manager (PM) is the way to combine several processes in a pipeline
- Sequential, parallel and conditional execution of processes
- The pipeline is launched according to DBMS state
- As the pipeline works, state is logged in the DBMS.
- When the pipeline completes, all processes terminate.

The Operating System: Problem or Solution?

- Services we required for this distributed system are no more available today than they were in 1987
- In fact, operating systems have stagnated for 20 years
- What have you guys been doing?
 - 10 Companies (avg)
 - x 10 Years (avg)
 - x \$1x10⁸ Software R&D (avg)
 -
 - = O(\$10¹⁰)
- This is real money