Phil Smith III

COMPUTER ASSOCIATES®

Agenda

- Background and history
- The Project
- Lessons learned
- Today and the future

Background

"E-mail is the killer app of the Internet"

— Jesse Berst, Ziff-Davis

Could you do your job without e-mail?

(Would you want to?)

COMPUTER® ASSOCIATES

Ford History

- Ford Motor Company founded in 1903
- Contrary to popular opinion, Henry Ford did ***not*** invent the automobile
  - ✓ He did perfect mass production manufacturing
- In 1915, over half of the automobiles sold in the U.S. were Fords!
- Pioneered higher wages, customer rebates

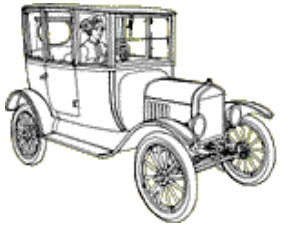

Ford Today

- The #2 U.S. manufacturer
- More than 6 million Fords built in 1999!
- 185,000 employees at plants in nine states, tens of thousands more overseas
- Annual IT spending about $1.5 billion (1% of total revenue, vs. 2.7 % for GM)
- Own Jaguar, Volvo, Aston-Martin, Hertz, part of Mazda VOLVO mazda LINCOLN Ford Mercury JAGUAR ASTON MARTIN
- Truly a world-class manufacturer

Ford and VM

- By any reckoning, Ford is a ***huge*** shop
  - ✓ Many MIPS of MVS, VSE, …and VM!
- Ford installed VM in the early 1980s, for PROFS
  - ✓ Of course, many other uses were found…
  - ✓ 130,000 OfficeVision users at peak (1997)
  - ✓ In 1993, Hewitt Associates said:

PROFS is the smartest, best, luckiest shot Ford ever took

- Six production images, most with 2GB main
  - ✓ Still over 30,000 unique VM logons daily

Ford Meets The Internet

- By 1996, Internet e-mail access need critical
- Project to find solution for OV users
  - ✓ Existing OV support deemed too weak
- Third-party product selected
  - ✓ SUN Solaris-based
  - ✓ Converts OV mail to SMTP and vice-versa
- Rolled out company-wide by 1997

Old Gateway Structure

- Four gateway machines:
  - ✓ OV ⇔ Ford Intranet (“FORDSMTP”)
  - ✓ OV ⇔ Exchange
  - ✓ OV ⇔ Internet
  - ✓ Test (development) machine
- Each server: 4–6 CPUs, 1.5GB RAM, 80 GB DASD …
- Also OS/2 boxes to do NJE ⇔ TCP/IP

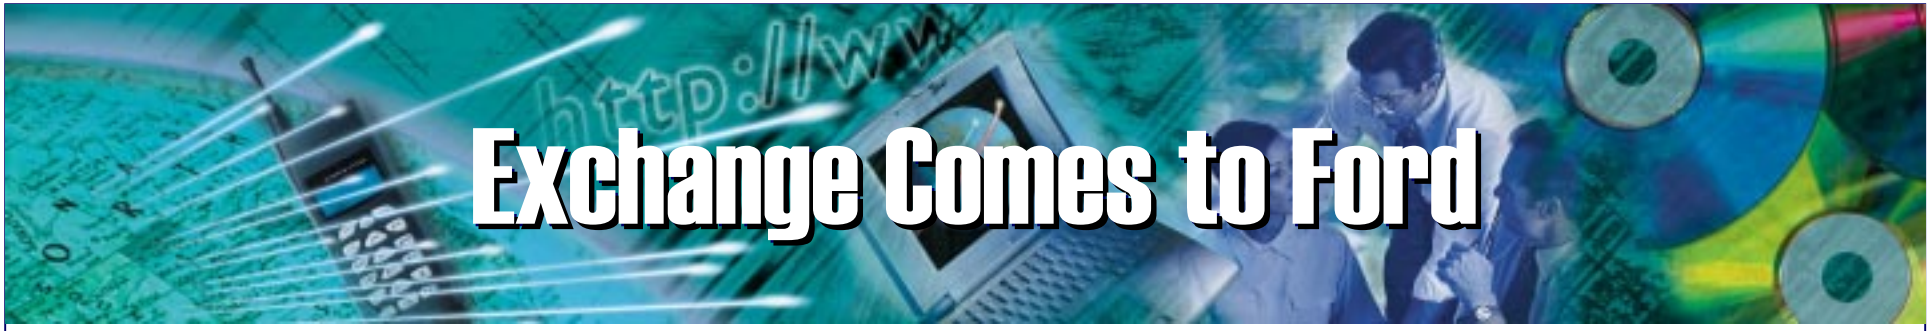

Exchange Comes to Ford

- In 1998, conversion from OV to Exchange started
  - ✓ Not without grumbling from users!
  - ✓ OV may be old, but it's fast, cheap, and people like it
- Today, more than half of OV users have been converted to Exchange, but ...
  - ✓ (A few) new OV users are still being added!
  - ✓ OV still the technology of choice for mass mailings

New Solution Needed

- Gateways overloaded, also not Y2K-ready
  - ✓ Delivery delays often 30 minutes
  - ✓ Nightly directory synchronization took hours
  - ✓ Migrated users had to wait for overnight synch
  - ✓ Server restarts required 45 minutes!
- Need for new solution identified

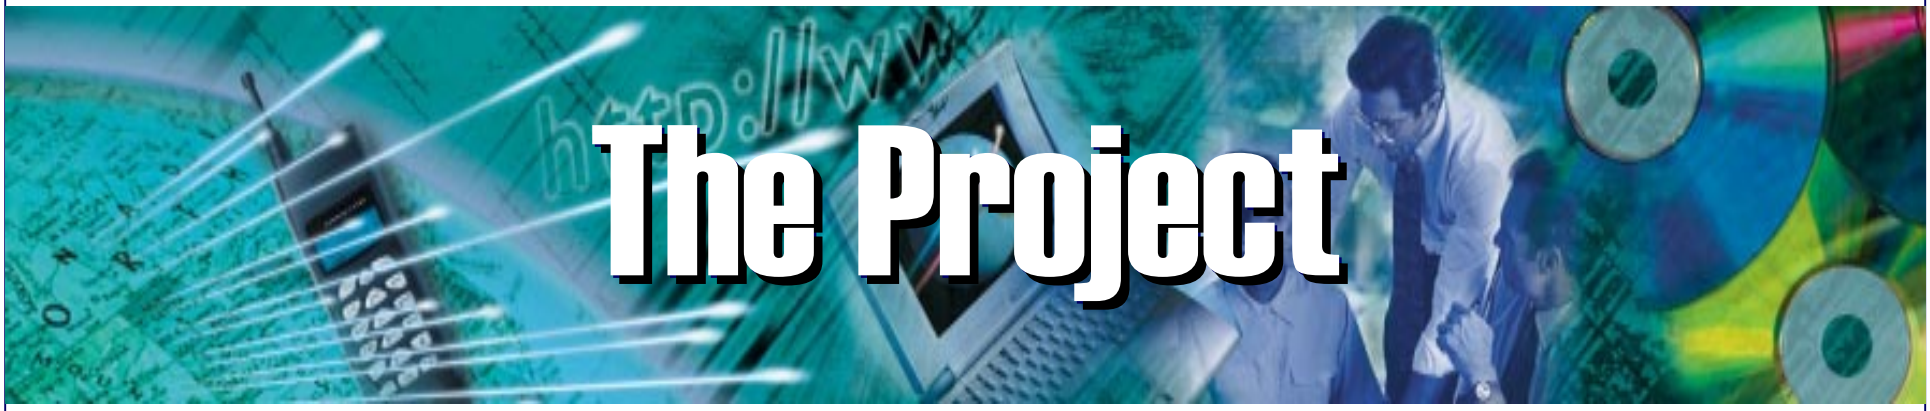

COMPUTER ASSOCIATES®

The Decision

- Put processing needed by VM users on VM
- No commercial products available
- IBM offered internal-use solution: XAgent
  - ✓ Heavily used for Internet-to-VNET e-mail
  - ✓ Collection of VM service virtual machines
  - ✓ Rexx, CMS Pipelines, **WAKEUP**
- Ford bought a copy of XAgent

The Approach

- XAgent not complete solution, much work required
- CA (Sterling, at the time) contracted to convert XAgent for Ford needs
- Six-month project, starting April 1999
- Al Lawrence & Vince Sheeran (Ford), me

The Results

- Some concepts retained, code entirely rewritten
  - ✓ New machines added for monitoring, statistics
- Much improved over XAgent
  - ✓ ***Not*** a criticism — we enjoyed several luxuries:
    - XAgent as a model
    - Example of how things should look to users
    - Time, commitment to meet Ford's specs
- 16KLoC of interpreted REXX
  - ✓ Compilation planned, but proved unnecessary

Rollout

- New gateway, dubbed **VM Connector**, rolled out in phases:
  - ✓ July 31: Exchange
  - ✓ August 28: Rest of Ford Intranet
  - ✓ September 26: Internet
  - ✓ Old gateways decommissioned September 30
- Constant minor enhancements since

How Did It Work?

- Users mostly failed to notice (this is a ***good*** thing!)
- Over 250,000 notes delivered ***per day***
- Latency reduced to (typically) sub-second
- Mailings of 10,000++ delivered in minutes

Connector Structure

- Mainline: **dealer** and **translator** machines
  - ✓ One or two dealers (MGDEIN, MGDEOUT)
  - ✓ One or more translators (MGTRIN*x*, MGTROUT*x)*
- Code is bidirectional
  - ✓ Machines can be unidirectional by virtue of what they're connected to (and thus receive)

- Dealer machines receive files and notes from RSCS and SMTP
- Decide inbound vs. outbound based on route, tag
- Forward to appropriate translator
- Files over specified size go to “big” translators (MGTRINB, MGTROUTB)
- Files over absolute maximum get bounced with message explaining why

Translators

- Translators receive and examine files
  - ✓ Decompose into address information, body
- Addresses resolved, addresses rewritten
- Meeting notices converted (both directions)
- Handles OV notes, SMTP mail, CMS notes

Translators

- Attachments for OV users detached
  - ✓ Translated as appropriate, per extension mapping table
  - ✓ Sent to reader in Netdata format
  - ✓ Users can **RECEIVE**, then download as binary
- Non-mail files MIME-encoded (for SMTP)
  - ✓ Enables CMS **SENDFILE** to LAN users

- Address resolution reflects long e-mail heritage at Ford: over 45 possible paths!
- Every Ford user has eight-character official userid ("**CDSid**")
- Ford Intranet users must register (acquire CDSids) to send through Connector
  - ✓ Mail from unregistered users gets bounced

Internet Addressing

- Internet senders get “autoregistered” as eight-character alias at node **EXTERNAL**
- OV users send to **EXTERNAL(*alias*)**
  - ✓ Unused aliases expire after 90 days
- Can also use “secondary addressing”:
  - ✓ Send to nickname **INTERNET**, aka userid **EXTERNAL(AAGENT)**
  - ✓ Put To:, Cc:, Bcc: lines in message body

Multiple Mailboxes

- Many users have OV ***and*** Exchange or other SMTP mail addresses
  - ✓ Gateway “knows” user’s preferred address, so all mail received in preferred mailbox
  - ✓ If user migrates to Exchange, no impact to senders (gateways updated within the hour)

- Address data comes from several sources:
  - ✓ SMTP file loaded at startup (~200,000 entries)
  - ✓ Autoregistration file, loaded at startup and updated dynamically (~400,000 entries)
  - ✓ CMS **RSLVNODE** utility maps CDSid to node (VM node, Exchange domain, FORDSMTP)
  - ✓ Startup latency: 70–90 seconds

"Bounced" Mail

- Undeliverable mail routed to "bouncer" machine (MGBOUNCE)
  - — Invalid recipient
  - — Unregistered sender
  - — File too large
  - — etc.
- Heuristics avoid attempted bounces to impossible SMTP addresses
- Copies of some bounces sent to Postmaster ID for later analysis

Virus Trapping

- Most PC viruses have a recognizable name
  - ✓ LOVE-LETTER-FOR-YOU.TXT.vbs, **`zipped_files.exe`**, **`pics4you.exe`**
- Translators process attachments
  - ✓ Recognize and delete known virus names
  - ✓ Add notice to message body about deleted attachment
- Not true virus scanning, but quite effective!
  - ✓ Especially for recent “LoveLetter” virus

Incoming Mail Flow

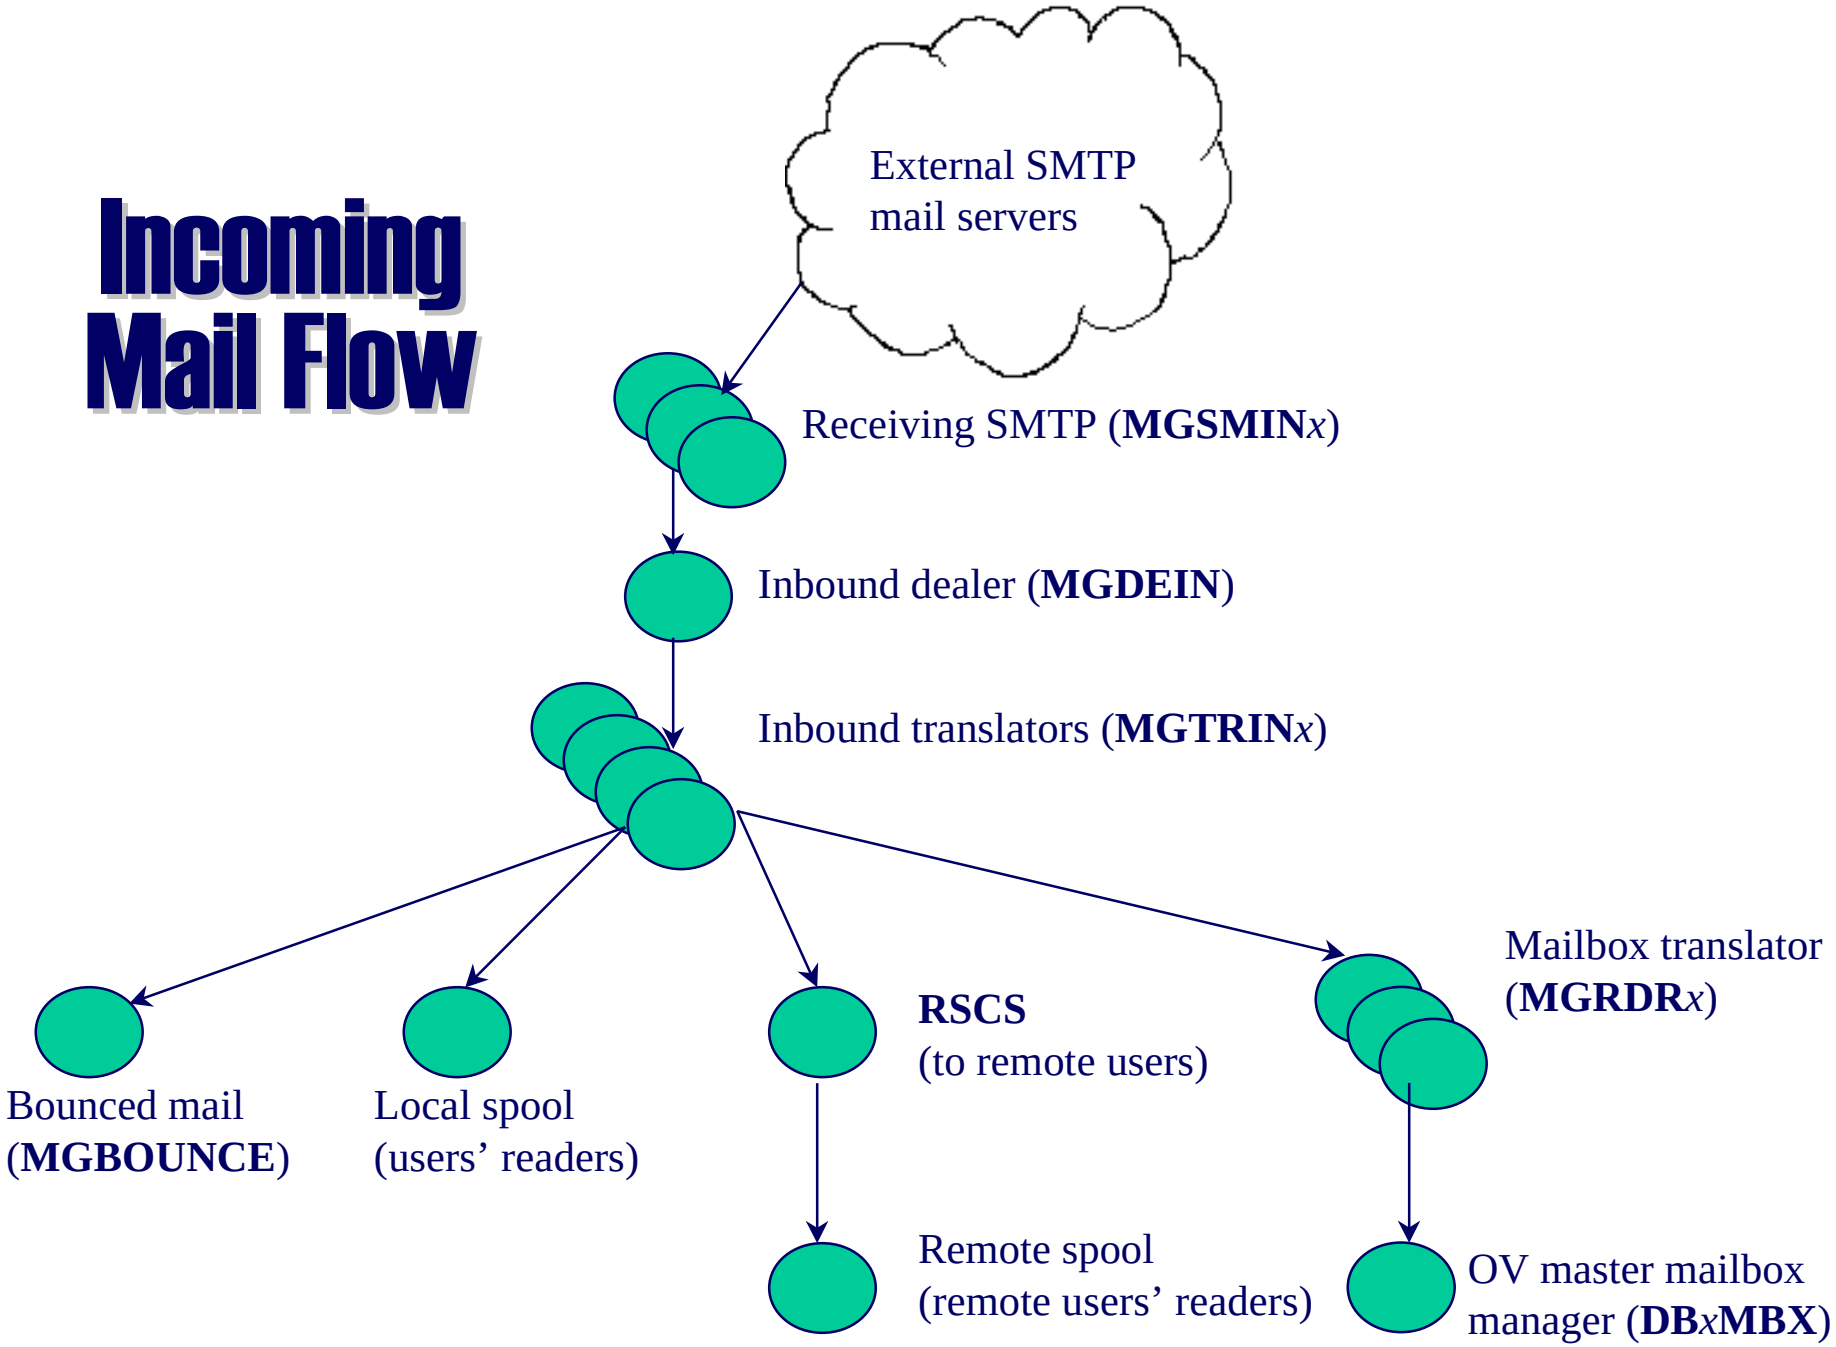

Outgoing Mail Flow

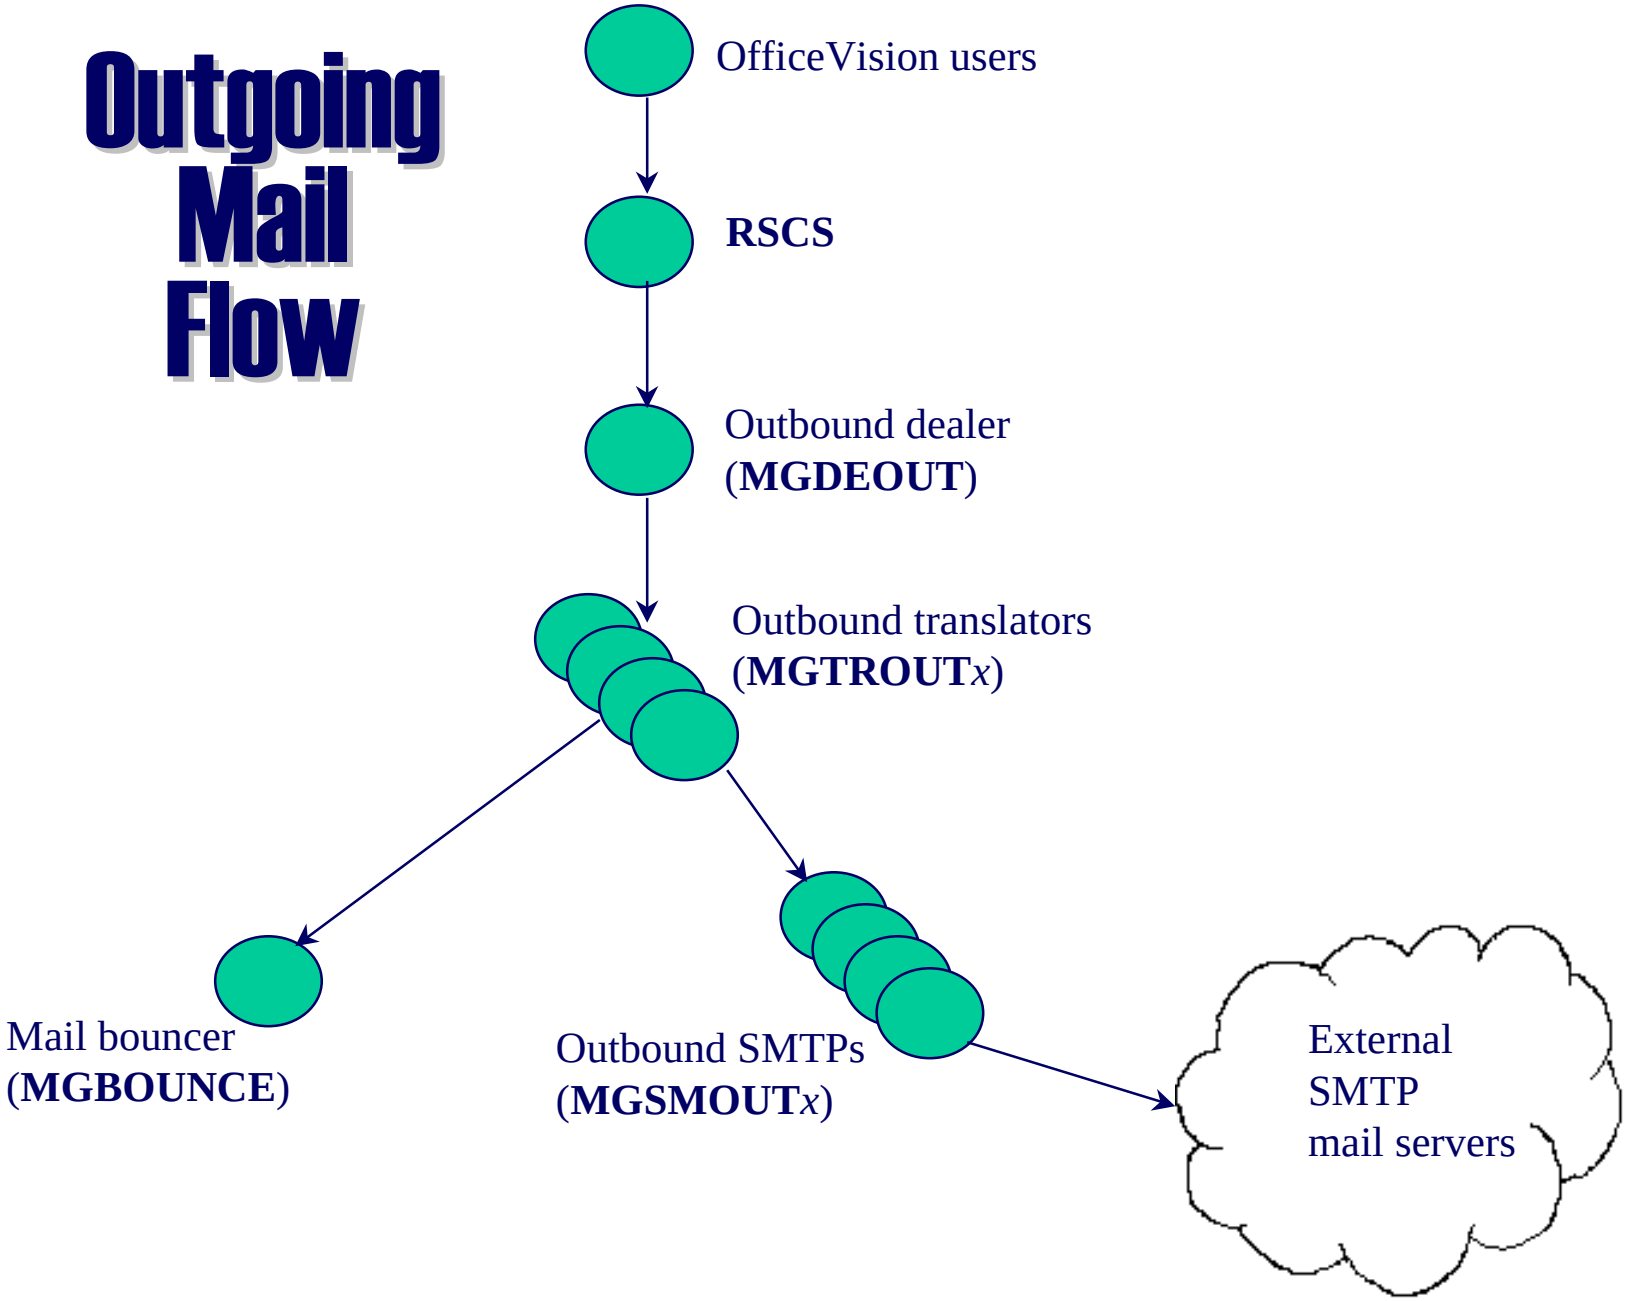

ReaderThief

- Original plan: use IBM OfficeVision ReaderThief PRPQ for OV delivery
- Bug in ReaderThief or CP causes *SPL errors reading files under heavy load
- Fix: restart ReaderThief periodically
  - ✓ Didn’t work — caused misdelivered mail

ReaderThief (continued)

- Next try: “round robin” files among ReaderThief machines
  - ✓ Caused ***more*** mis-delivered mail!
- Al spent Labor Day analyzing ReaderThief
  - ✓ Reduced needed function to single ReaderThief Pipes stage
  - ✓ Wrote MGREADER to use that stage
- No problems since!

Dealer Switching

- Dealer machines examine translator queues every 20 files
- Switch translators if current has backlog
- Provides automatic capacity indicator:
  - ✓ If all machines under heavy use, more needed
  - ✓ If “last” machine in sequence mostly idle, capacity OK

Statistics and Monitoring

- MGWATCH machine “watches” all other servers, checks health every 10 seconds
  - ✓ Intranet Web page shows gateway status (mainly for Helpdesk)
- Statistics generated show file counts, CPU time, latency, etc., etc…
  - ✓ MGSTATS produces reports

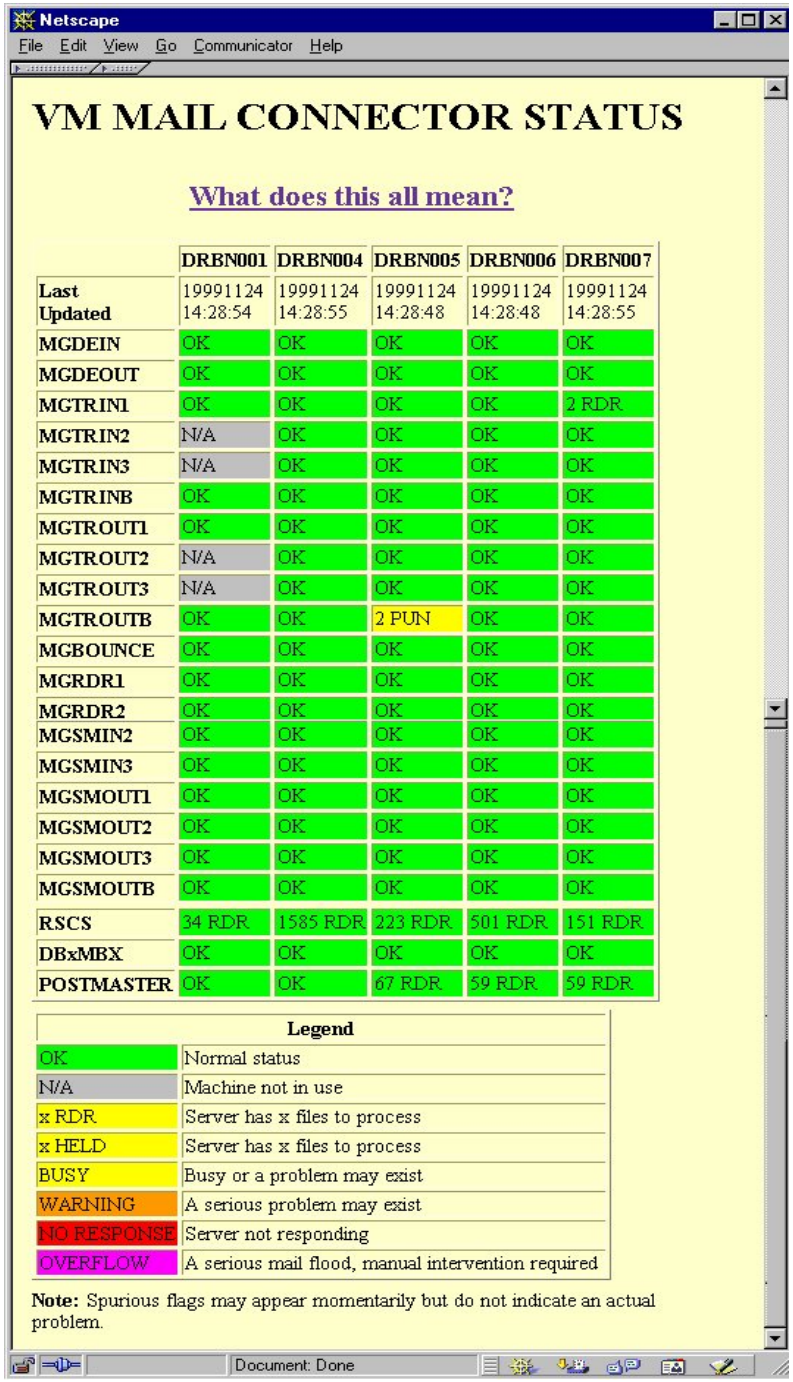

	DRBN001	DRBN004	DRBN005	DRBN006	DRBN007
Last Updated	19991124 14:28:54	19991124 14:28:55	19991124 14:28:48	19991124 14:28:48	19991124 14:28:55
MGDEIN	OK	OK	OK	OK	OK
MGDEOUT	OK	OK	OK	OK	OK
MGTRIN1	OK	OK	OK	OK	2 RDR
MGTRIN2	N/A	OK	OK	OK	OK
MGTRIN3	N/A	OK	OK	OK	OK
MGTRINB	OK	OK	OK	OK	OK
MGTROUT1	OK	OK	OK	OK	OK
MGTROUT2	N/A	OK	OK	OK	OK
MGTROUT3	N/A	OK	OK	OK	OK
MGTROUTB	OK	OK	2 PUN	OK	OK
MGBOUNCE	OK	OK	OK	OK	OK
MGRDR1	OK	OK	OK	OK	OK
MGRDR2	OK	OK	OK	OK	OK
MGSMIN2	OK	OK	OK	OK	OK
MGSMIN3	OK	OK	OK	OK	OK
MGSMOUT1	OK	OK	OK	OK	OK
MGSMOUT2	OK	OK	OK	OK	OK
MGSMOUT3	OK	OK	OK	OK	OK
MGSMOUTB	OK	OK	OK	OK	OK
RSCS	34 RDR	1585 RDR	223 RDR	501 RDR	151 RDR
DBxMBX	OK	OK	OK	OK	OK
POSTMASTER	OK	OK	67 RDR	59 RDR	59 RDR

Legend	
OK	Normal status
N/A	Machine not in use
x RDR	Server has x files to process
x HELD	Server has x files to process
BUSY	Busy or a problem may exist
WARNING	A serious problem may exist
NO RESPONSE	Server not responding
OVERFLOW	A serious mail flood, manual intervention required

Note: Spurious flags may appear momentarily but do not indicate an actual problem.

Reliability

- Machines can (mostly) fend for themselves:
  - ✓ Rexx syntax errors logged, MSGs/pages sent
  - ✓ Restart automatically at top of main loop (avoids startup latency)
- Errors decoding MIME cause note copies to Postmaster, undecoded notes to recipient
  - ✓ Allowed detecting, fixing many strange cases

- Built-in debugging facilities help:
  - ✓ **Debug** settings set Rexx trace points
  - ✓ **DebugTo** sends copies of input, output files
  - ✓ **DebugDeal** deals files from selected senders to specific translators
  - ✓ **REXX** command allows querying/setting internal variables on-the-fly

COMPUTER ASSOCIATES®

A Pipes Performance Lesson

- CMS Pipelines are remarkably powerful
  (we knew that)
- Rexx variable interface is remarkably slow
  (we knew that…didn't realize ***how*** slow!)
  - ✓ E-mail bodies originally kept in variables
  - ✓ Large messages took up to 30 minutes (ouch)
  - ✓ Changed to use disk files: now max 30 seconds

Another Performance Lesson

- Dealers lagged translators during mass mailings
- Seemed unlikely, since dealers do very little
- Noticed: once backlog below 100, dealers ***flew***
- Bingo: 8192 = default Pipes CP stage buffer, approximately 100 files' worth of output
  - ✓ Retries QUERY with larger buffer until output fits
- Specifying explicit CP stage buffer size fixed

Standards Aren't

- MIME-type for UUENCODEd attachments:

  `Content-Type: X-UUENCODE`

  `Content-Type: UUENCODE`

  `Content-Type: X-UUE`

- Even standard MIME types can be bizarre:

  `Content-Type: X-Zm-BASE64`

  - ✓ Just **BASE64**, not sure why weird type

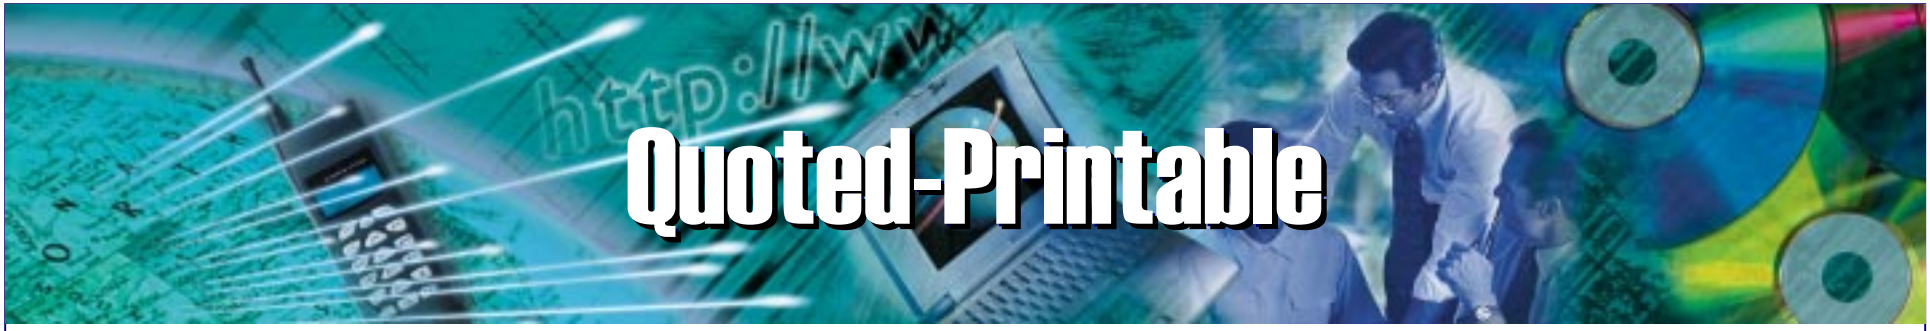

Quoted-Printable

- Replaces “dangerous” characters with “***=xx***”
- Quoted-printable notes often aren’t:
  - ✓ “Stalker’s Mailer” QP-encodes headers!
  - ✓ Other mailers QP-encode, ***then*** append .sig (many sigs have equals signs in them!)
- Solution: when invalid QP sequence found, ignore, keep scanning
- John Hartmann provided QPDECODE stage

Looping Mail Happens

- Early on, note somehow addressed ***to*** SMTP made several hundred thousand trips
- Unregistered Ford internal servers exist
  - ✓ Send notes (alerts, etc.)
  - ✓ Connector sends bounce
  - ✓ Server sends response saying “Huh?”
  - ✓ Connector sends bounce ...
- We’re now ***really*** good at detecting loops!

Unregistered Automated Servers

- There are ***many*** automated servers out there
  - ✓ Some send hundreds of alerts ***per hour***
  - ✓ Amazingly, some of these persist for weeks, without owners noticing lack of delivery
  - ✓ When owner cannot be identified, bounces continue until server's mailbox fills…
- **NOBOUNCE** list suppresses bounce to ill-behaved servers, avoids loops

Meeting Notice Quirks

- External recipients get human-readable notices
  - ✓ But: date is in ***sender's*** local Windows format
  - ✓ Mostly OK in 1999, but when is 01/02/03?
- Forwarded notices show ***creator*** as origin!
  - ✓ Bad if reply has comments about him/her…!
- External senders use vCalendar:
  - ✓ Outlook 98 only handles vCal V1.0, ignores recurrence
  - ✓ Outlook 2000 handles vCal V2.0, with recurrence
  - ✓ Solution: create both formats, add verbiage to note telling user to select appropriate attachment

- IBMMail sends ***lots*** of funky notes:
  - ✓ Pseudo-OfficeVision format, with x‘FE’ first line, but RFC822-style headers
- Massive amounts of SPAM
  - ✓ Fortunately, part of project was to eliminate costly IBMMail connection…
    … kind of solved SPAM problem!
- Did have to handle strange formats in interim

Shared File System Issues

- CMS Shared File System powerful, but locking problematic:
  - ✓ No native “enqueue on lock” (can loop, but…)
  - ✓ Locking by ***userid***, not userid and node
- Heavy lock use by same userids on multiple systems confuses SFS
  - ✓ Seems worse when DIAGNOSE x‘D4’ in use
  - ✓ Unresolved: avoided through enqueue file
- Could use CMS Multitasking thread services

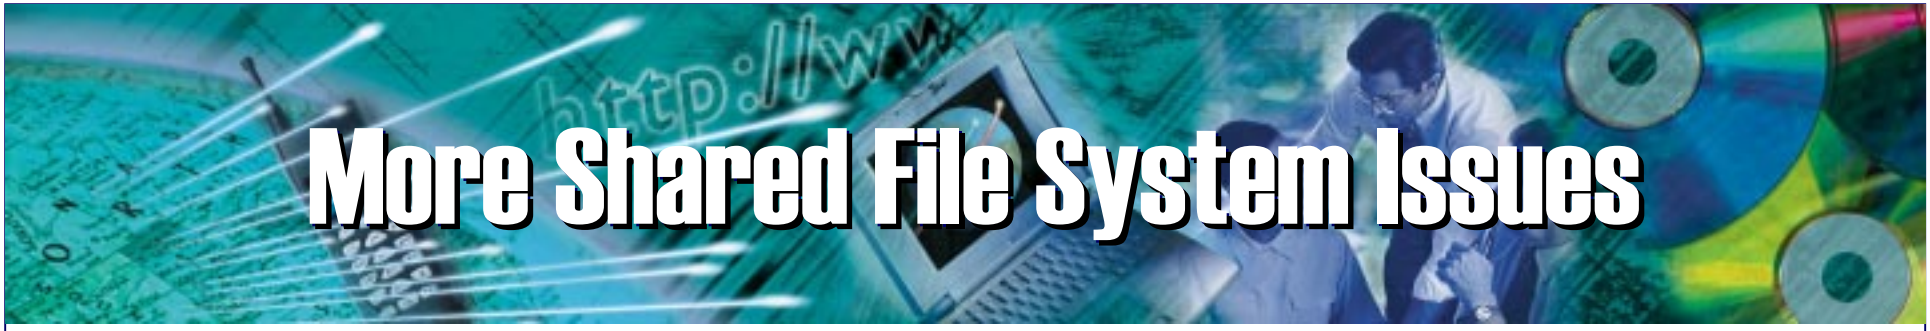

More Shared File System Issues

- Early on, SFS bug caused server hangs
- Seemed to be related to heavy use
- After report, tracing, etc., IBM APAR VM62301 fixed
- A must for SFS-intensive shops!

Other Issues

- If multiple virtual readers defined, **`WAKEUP (RDR`** results seem wrong
  - ✓ Actually they make sense, but are unintuitive; see APAR VM62207
- SMTP sends "**`Note delivered`**" msg
  - ✓ Ford users ***strongly*** disliked this
  - ✓ IBM provided customized SMTP to suppress
  - ✓ Will be configuration option in future release

Drowning In Mail

- Ford is an extremely dynamic shop
- SFS-related hangs caused painful backlogs:
  - ✓ VM SPOOL file limit of 9999/user reached in less than an hour
  - ✓ Then RSCS and SMTP clogged up…
  - ✓ Once fixed, backlogs took time to clear
- Would be less painful with current performance improvements

Fixing Problems

- All machines run with
  `CP TERMINAL MORE 0 0`
  `CP TERMINAL HOLD OFF`
- Cannot even disconnect, if connected to debug problem and **`MORE`** not reset!
  - ✓ If mass mailing comes through while connected, forget about it…

COMPUTER® ASSOCIATES

Current Status

- VM Connector in full production at Ford
  - ✓ Users are happy
  - ✓ Management is happy
- Over $1M savings projected in FY2000
- SUN & OS/2 hardware freed for other uses
- Staff load reduced
  - ✓ Full-time gateway admin no longer needed
  - ✓ No new VM staff required

Conclusion

- A true success story for VM!
- VM ***veni, vidi, vici!***
- Workstation hardware and software replaced, results ***much*** improved!

There are other opportunities!

People Who Helped

- WAKEUP support: Colleen Brown (IBM)
- SFS: Jim Wallace (IBM)
- VM TCP/IP: Glenn Skryp, Romney White (IBM)
- OfficeVision: Tracy Dean (IBM)
- XAgent: Dave Martin, Larry Nomer (IBM)
- The Piper, John Hartmann (IBM)
- Melinda Varian, for Pipes help (Princeton)

Contact Info

Phil Smith III
Computer Associates International, Inc.
1800 Alexander Bell Drive
Reston, VA 20191

(703) 264-8514 (voice)
(703) 264-8190 (FAX)

Phil.Smith@Sterling.com
USSCIPHS at IBMMail
VMSysprog on AIM