The CMS Shared File System: Usage and Administration

Daniel P. Martin
Rocket Software, Inc.

Monday, August 10th, 2015
Session Number 17274

#SHAREorg

Agenda

- Intro: Where is this person from, and why is he here?
- Background: What's a File System?
- History Lesson: CMS File Systems, Then & Now
- SFS: It's already inside the house…
- SFS: Taking charge, and looking under the hood
- SFS: Care and Feeding of a File Pool Server
- SFS: End User Interaction
- SFS: A Quick Look at BFS, OpenExtensions and POSIX
- Summary
- Additional Resources

Introduction
About Rocket Software

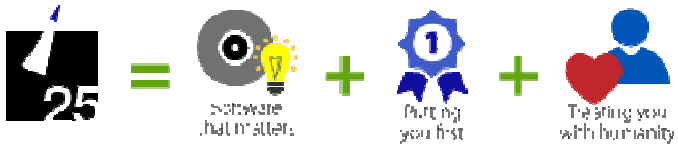

http://www.rocketsoftware.com/about

We Build Software That Matters.

That's not rocket science. That's Rocket Software.

One day in 1990, one of our founders started writing assembler code in a spare bedroom in his house near Boston. The very first product he built helped large enterprises solve an important problem—how to process more database queries more efficiently. That was how Rocket Software started.

More than two decades later, our software engineers have built well over 100 products that solve problems across a broad spectrum of enterprise technology. Our engineers talk with you, our customers and partners, every day to discover new pain points and learn about the (hopefully really hard) problems and challenges you face.

We help you prevent outages, protect your data, store your data, share your data, virtualize your data, manage your networks, improve your service levels, discover insights, modernize your applications, access and connect users and applications, minimize risk and increase compliance, and so much more. We build and deliver products that matter to you so that you can deliver your best products, solutions, and services to your customers and grow your business.

We Put You First.

Rocket was founded on the premise that we would build products that matter for people—and we have never wavered from that. In fact, it's the cornerstone of our core values that we live every day—we put you, our customers and partners, first.

We treat all of our customers and partners as individuals rather than transactions. That's been our history. And that's why our customers and partners see us as a trusted partner.

We don't just sell software. We care about our customers' and partners' success—a win for you is a win for us. We spend our time solving the problems that keep you awake. We build software that matters—to you.

Treating You With Humanity.

This is the piece that almost every business today gets wrong. This is the piece that we talk about getting right at Rocket. We put you first and are committed to never letting you fail. We are committed to each other. Rocketeers don't let each other fail. We are people building software for people. We are people solving problems for people. We have worked really hard over the past 25 years to earn your trust. And for the past 25 years we have always tried to treat you with humanity.

About your speaker

- 23 Years as a Customer
  - 22 years at University of Arkansas
    - …starting with VM/370 R6 PLC 3 (CDC Omega 480-III, anyone?)
    - …lots of Unix / Linux (Solaris, HP-UX, AIX, NCR/Teradata, Various distros)
    - …IT Security Weasel*
      - *"Telling people things they didn't want to hear before it was cool."
  - A brief detour through "A Certain Major Retailer"**
    - **It was the early 1980's – We called it "Distributed Systems"
- 12-plus years and counting: Senior Software Developer for Rocket Software since 2003
  - Lead design / development
    - IBM Backup and Restore Manager for z/VM
    - IBM Archive Manager for z/VM
  - Co-conspirator
    - IBM Tape Manager for z/VM
    - IBM Operations Manager for z/VM
- Away From the Keyboard:
  - Certified Law Enforcement Officer, Search and Rescue Worker, Emergency Medical First Responder, and misplaced farm boy with an idiosyncratic fondness for shiny things…

Target Audience

- Some familiarity with IBM z Systems Concepts, Facilities, and Terminology
- Some familiarity with IBM z/VM Concepts and Terminology
- Foundation-level Introduction to the Topic

Background: What's a file system?

Background: What's a file system?

Proposition:

A file system organizes raw data…

Background: What's a file system?

…into structured information…

Background: What's a file system?

…and makes it available in a consistent form…

Background: What's a file system?

… for use by applications.

Example: Here's one file, out of a file system which contains thousands of items.

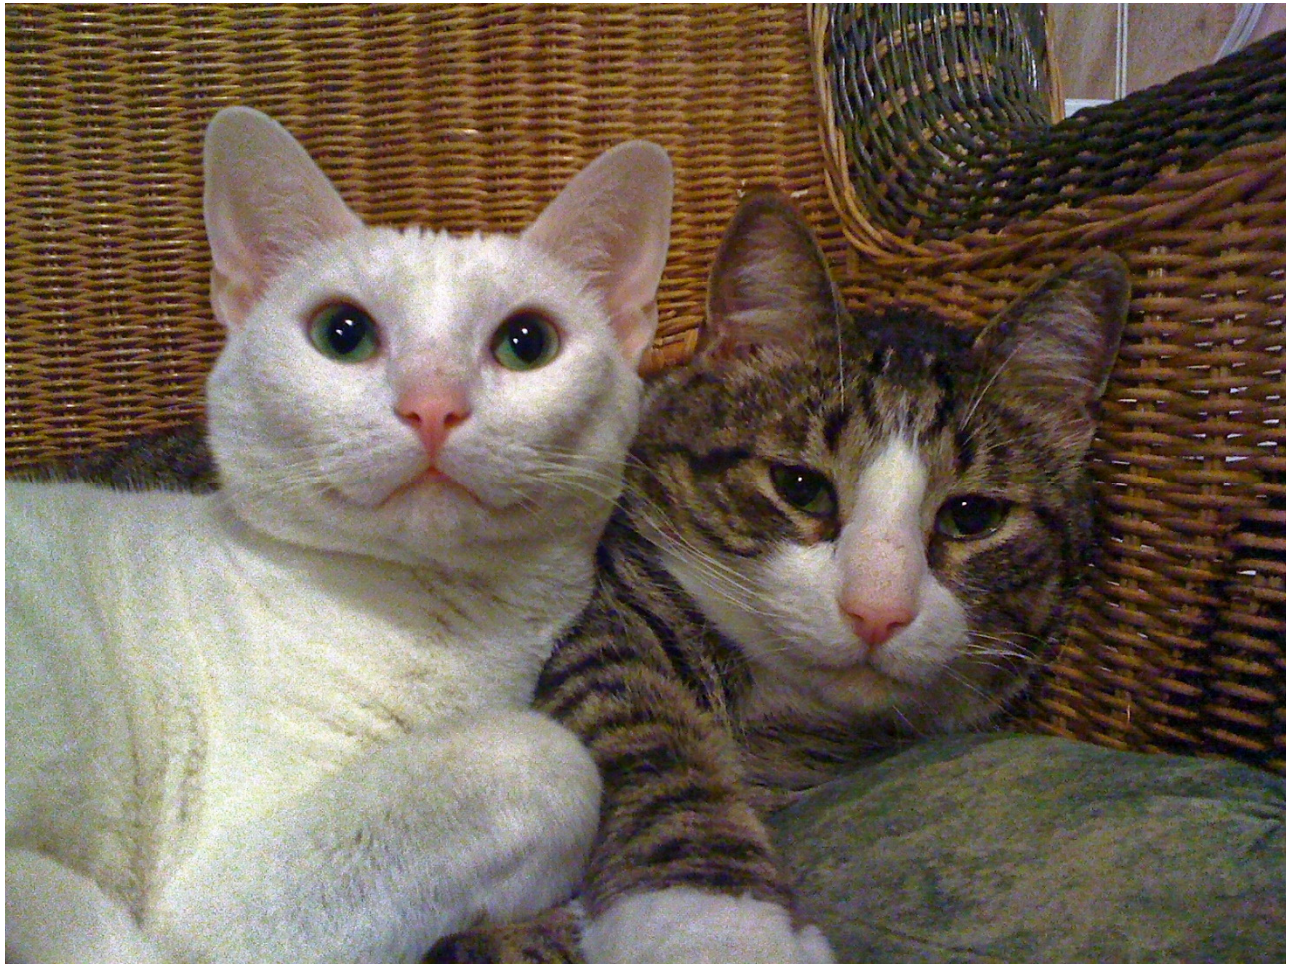

Example: Here's one file, out of a file system which contains thousands of items.

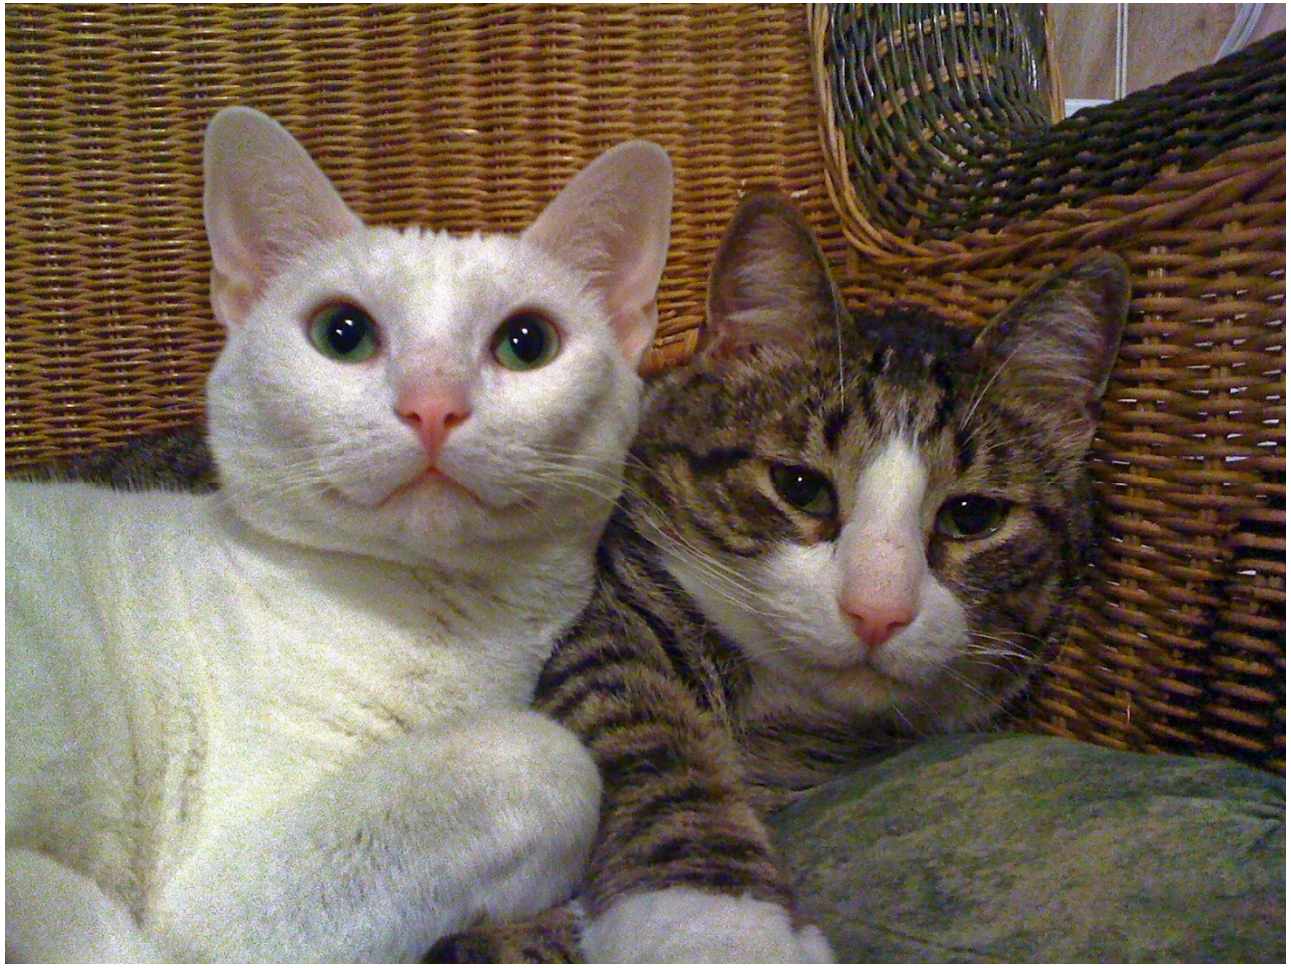

The 5,000 year trek toward modern civilization has made it possible for you to watch cat videos on your smart phone while I'm talking.

Example: Here's one file, out of a file system which contains thousands of items.

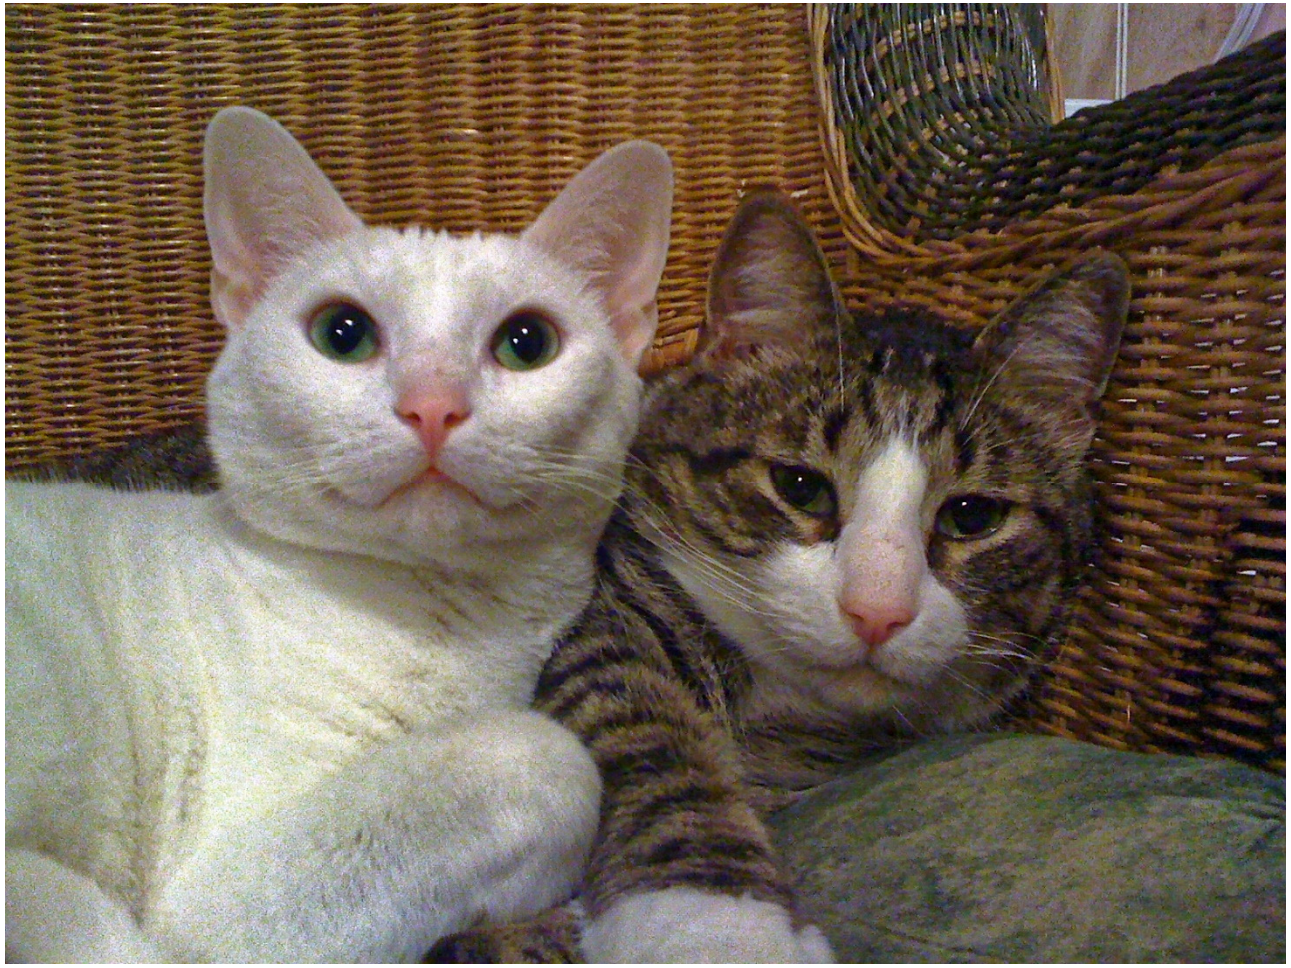

The 5,000 year trek toward modern civilization has made it possible for you to watch cat videos on your smart phone while I'm talking.

We call this "Progress."

You're right.

You're right.
That was a ridiculous example.

…but think about all of the steps it took to get from "there and then" to "here and now."*

* The cats weren't impressed, either.

You're right.
That was a ridiculous example…

…because the history behind the technology is worth a look.

z/VM History Lesson:
CMS File Systems, Then and Now

- In the beginning, there were Virtual Machines…

z/VM History Lesson:
CMS File Systems, Then and Now

- In the beginning, there were Virtual Machines…
- …and VM used minidisks, just like now.

z/VM History Lesson:
CMS File Systems, Then and Now

- In the beginning, there were Virtual Machines…
- …and the only CMS file system was minidisk-based.

z/VM History Lesson:
CMS File Systems, Then and Now

- …and the only CMS file system was minidisk-based.
- **Point of reference**: In 1973, the new IBM 3330-11 could store 200MB on a 25-pound disk pack.
  - (That's about 7.1 pounds per average-size MP3 file…)
  - …which means that your 32GB smart phone would weigh almost 25 tons.*

*batteries not included.

z/VM History Lesson: CMS File Systems, Then and Now

- Some things have changed since then:
  - Storage density, reliability, and performance have increased.
  - Storage costs have declined.

z/VM History Lesson: CMS File Systems, Then and Now

- Other things haven't changed… much:
  - The z/VM hypervisor (CP) "just" virtualizes hardware architecture.
  - CP isn't (much*) aware of virtual machine DASD I/O, as long as it complies with the architecture.
  - The CMS operating system still uses a minidisk-based file system (EDF – Enhanced Disk Format).
    - *CP *does* know how to read EDF file systems; that's where the operating system and configuration files load from now. We have some lab sessions for that.

z/VM History Lesson: CMS File Systems, Then and Now

- CMS EDF* minidisk file systems
  - **GREAT** for read-only sharing
    - All the activity your I/O subsystem will carry
      - …as long as nobody updates the file system.

*EDF: Enhanced Disk Format

…because CDF needed enhancement, but that's another story.

z/VM History Lesson: CMS File Systems, Then and Now

- CMS EDF minidisk file systems
  - **NOT** at all great for multiple-write access sharing
    - No serialization mechanism *
    - **"Last writer wins."**
      - If, by "win," you mean "mutually assured destruction."
      - * CP, the hypervisor, does support reserve/release because it's part of the architecture. CMS does *not* use reserve/release.

z/VM History Lesson: CMS File Systems, Then and Now

- CMS EDF minidisk file systems
  - Not so good for “one writer / many readers” sharing:
    - No handshaking with other CMS virtual machines.
    - I/O Caching by CP and DASD subsystems adds another layer of complication.
      - Cache is a wonderful thing, but the axiom “There’s no I/O like no I/O” will only carry you just so far…

z/VM History Lesson: CMS File Systems, Then and Now

- CMS EDF minidisk file systems

- Summary:
  - The CMS minidisk file system was never designed to share well with others.
  - CP is “only” responsible for ensuring architectural compliance.
  - Shared DASD environments, whether virtual or real, are a pain.
    - CMS is “just one more guest operating system” as far as CP is concerned.
    - Shared minidisks are one more form of shared DASD.

Problem:
We need a new file system for CMS

- Requirements:
  - Concurrent read-write access by multiple CMS guests.
  - Consistent file view from OPEN to CLOSE
  - Transaction level recovery (commit / rollback)
  - Serialization (lock / unlock)
  - Security:
    - CMS guests must not have access to DASD
    - Must provide Access control & Audit mechanisms
  - **MUST be backward-compatible with existing CMS applications**
  - While you're at it, there's this POSIX thing…

Q: How do we do this?
A: It's VM; put it in a service virtual machine.

- Basic moving parts of a File Pool Server:
  - Some EDF-format minidisks
    - Special "RESERVE" format
    - Some help from CP: *BLOCKIO service
  - Some configuration files
    - Resource enumeration
    - Privileged clients
    - Tuning parameters
  - Client/Server data transfer mechanism*
    - APPC (Application Peer-to-Peer Communication)
      - *VM: Client/Server before Client/Server was Cool

Q: How do we do this?
A: It's VM; put it in a service virtual machine.

- Compartmentalization and Containment
  - Client virtual machine "I/O" is via APPC
  - …with NO direct access to File Pool Server media
  - …with "good enough" access control baked in
  - …with standard, auditable ESM (i.e. RACF) as an option
  - …with a secure transport layer provided by the hypervisor

Q: How do we do this?
A: It's VM; put it in a service virtual machine.

- Compartmentalization and Containment:
  - File mode 0? That's not security.
  - This is not a new problem: access to storage media means access to everything stored on the media.
  - Encryption is good. I approve of Encryption.
  - Encryption is, at best, an arms race instead of a guarantee of security. All I need is time and horsepower.

Who did the heavy lifting?

- IBM VM Development
- Mid-1980's
  - …so "VM," not "z/VM"…
- Generally Available in VM/SP Release 6, November 1988
- The initial deployment had a few rough edges

The customer community generously responded with lots of well-reasoned, insightful feedback.

Q: Why do old-timers look askance at SFS?
A: Because… they're old…?

- Consider the state of the industry, around second half of the 1980's:

State of the Industry, second half of the 1980's:

Cover Page from the June 15th, 1985 edition of DATAMATION:

- "Desktop Computing" is just becoming a "thing."

- One might reasonably assume work on SFS was in progress around this time, but I couldn't possibly comment.

State of the Industry, second half of the 1980's:

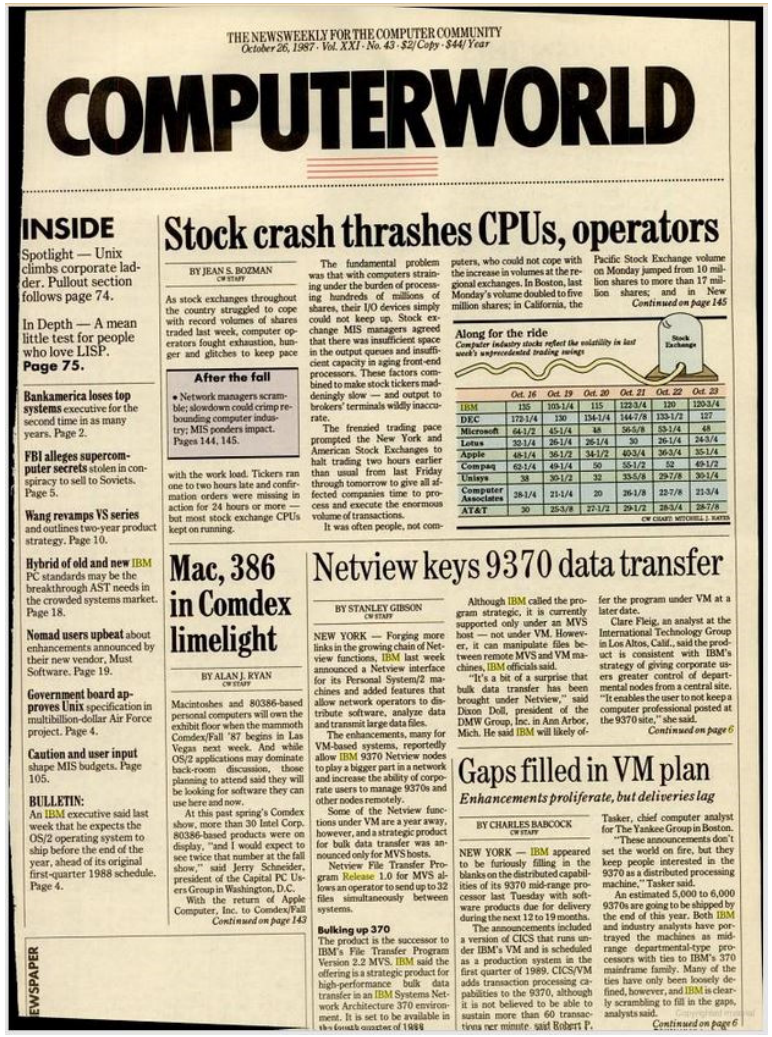

THE NEWSWEEKLY FOR THE COMPUTER COMMUNITY
October 26, 1987 · Vol. XXI · No. 43 · $2/Copy · $44/Year

COMPUTERWORLD

INSIDE

Spotlight — Unix climbs corporate ladder. Pullout section follows page 74.

In Depth — A mean little test for people who love LISP. Page 75.

Bankamerica loses top systems executive for the second time in as many years. Page 2.

FBI alleges supercomputer secrets stolen in conspiracy to sell to Soviets. Page 5.

Wang revamps VS series and outlines two-year product strategy. Page 10.

Hybrid of old and new IBM PC standards may be the breakthrough AST needs in the crowded systems market. Page 18.

Nomad users upbeat about enhancements announced by their new vendor, Must Software. Page 19.

Government board approves Unix specification in multibillion-dollar Air Force project. Page 4.

Caution and user input shape MIS budgets. Page 105.

BULLETIN:
An IBM executive said last week that he expects the OS/2 operating system to ship before the end of the year, ahead of its original first-quarter 1988 schedule. Page 4.

Stock crash thrashes CPUs, operators

BY JEAN S. BOZMAN
CW STAFF

As stock exchanges throughout the country struggled to cope with record volumes of shares traded last week, computer operators fought exhaustion, hunger and glitches to keep pace with the work load. Tickers ran one to two hours late and confirmation orders were missing in action for 24 hours or more — but most stock exchange CPUs kept on running.

After the fall

- Network managers scramble; slowdown could crimp rebounding computer industry; MIS ponders impact. Pages 144, 145.

The fundamental problem was that with computers straining under the burden of processing hundreds of millions of shares, their I/O devices simply could not keep up. Stock exchange MIS managers agreed that there was insufficient space in the output queues and insufficient capacity in aging front-end processors. These factors combined to make stock tickers maddeningly slow — and output to brokers' terminals wildly inaccurate.

The frenzied trading pace prompted the New York and American Stock Exchanges to halt trading two hours earlier than usual from last Friday through tomorrow to give all affected companies time to process and execute the enormous volume of transactions.

It was often people, not computers, who could not cope with the increase in volumes at the regional exchanges. In Boston, last Monday's volume doubled to five million shares; in California, the Pacific Stock Exchange volume on Monday jumped from 10 million shares to more than 17 million shares; and in New

Continued on page 145

Along for the ride
Computer industry stocks reflect the volatility in last week's unprecedented trading swings

	Oct. 16	Oct. 19	Oct. 20	Oct. 21	Oct. 22	Oct. 23
IBM	135	103-1/4	115	122-3/4	120	120-3/4
DEC	173-1/4	130	134-1/4	144-7/8	133-1/2	127
Microsoft	64-1/2	45-1/4	48	56-5/8	53-1/4	48
Lotus	33-1/4	26-1/4	26-1/4	30	26-1/4	24-3/4
Apple	48-1/4	36-1/2	34-1/2	40-3/4	36-3/4	35-1/4
Compaq	62-1/4	49-1/4	50	55-1/2	52	49-1/2
Unisys	38	30-1/2	32	35-5/8	29-7/8	30-1/4
Computer Associates	28-1/4	21-1/4	20	26-1/8	22-7/8	21-3/4
AT&T	30	25-3/8	27-1/2	29-1/2	28-3/4	28-7/8

CW CHART: MITCHELL J. KATES

Mac, 386 in Comdex limelight

BY ALAN J. RYAN
CW STAFF

Macintoshes and 80386-based personal computers will own the exhibit floor when the mammoth Comdex/Fall '87 begins in Las Vegas next week. And while OS/2 applications may dominate back-room discussion, those planning to attend said they will be looking for software they can use here and now.

At this past spring's Comdex show, more than 30 Intel Corp. 80386-based products were on display, "and I would expect to see twice that number at the fall show," said Jerry Schneider, president of the Capital PC Users Group in Washington, D.C.

With the return of Apple Computer, Inc. to Comdex/Fall

Continued on page 143

Netview keys 9370 data transfer

BY STANLEY GIBSON
CW STAFF

NEW YORK — Forging more links in the growing chain of Netview functions, IBM last week announced a Netview interface for its Personal System/2 machines and added features that allow network operators to distribute software, analyze data and transmit large data files.

The enhancements, many for VM-based systems, reportedly allow IBM 9370 Netview nodes to play a bigger part in a network and increase the ability of corporate users to manage 9370s and other nodes remotely.

Some of the Netview functions under VM are a year away, however, and a strategic product for bulk data transfer was announced only for MVS hosts.

Netview File Transfer Program Release 1.0 for MVS allows an operator to send up to 32 files simultaneously between systems.

Bulking up 370
The product is the successor to IBM's File Transfer Program Version 2.2 MVS. IBM said the offering is a strategic product for high-performance bulk data transfer in an IBM Systems Network Architecture 370 environment. It is set to be available in the fourth quarter of 1988.

Although IBM called the program strategic, it is currently supported only under an MVS host — not under VM. However, it can manipulate files between remote MVS and VM machines, IBM officials said.

"It's a bit of a surprise that bulk data transfer has been brought under Netview," said Dixon Doll, president of the DMW Group, Inc. in Ann Arbor, Mich. He said IBM will likely offer the program under VM at a later date.

Clare Fleig, an analyst at the International Technology Group in Los Altos, Calif., said the product is consistent with IBM's strategy of giving corporate users greater control of departmental nodes from a central site. "It enables the user to not keep a computer professional posted at the 9370 site," she said.

Continued on page 6

Gaps filled in VM plan

Enhancements proliferate, but deliveries lag

BY CHARLES BABCOCK
CW STAFF

NEW YORK — IBM appeared to be furiously filling in the blanks on the distributed capabilities of its 9370 mid-range processor last Tuesday with software products due for delivery during the next 12 to 19 months.

The announcements included a version of CICS that runs under IBM's VM and is scheduled as a production system in the first quarter of 1989. CICS/VM adds transaction processing capabilities to the 9370, although it is not believed to be able to sustain more than 60 transactions per minute, said Robert P. Tasker, chief computer analyst for The Yankee Group in Boston.

"These announcements don't set the world on fire, but they keep people interested in the 9370 as a distributed processing machine," Tasker said.

An estimated 5,000 to 6,000 9370s are going to be shipped by the end of this year. Both IBM and industry analysts have portrayed the machines as mid-range departmental-type processors with ties to IBM's 370 mainframe family. Many of the ties have only been loosely defined, however, and IBM is clearly scrambling to fill in the gaps, analysts said.

Continued on page 6

Cover Page from the October 26th, 1987 edition of COMPUTERWORLD:

- The Macintosh, the intel 80386 PC, and OS/2 are hot topics at Comdex.
- Netview. 9370's. The PS/2.
- "**Gaps filled in VM plan**" CICS/VM? (Yes. It was a real thing.)

State of the Industry, second half of the 1980's:

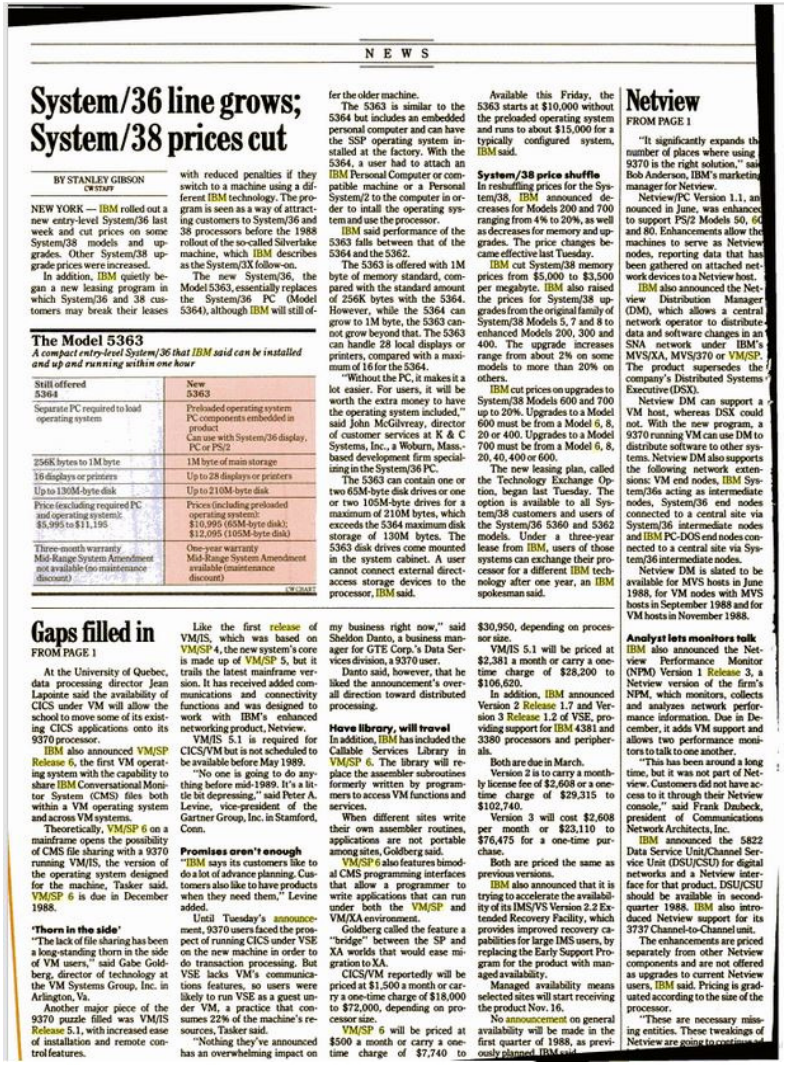

NEWS

System/36 line grows; System/38 prices cut

BY STANLEY GIBSON
CW STAFF

NEW YORK — IBM rolled out a new entry-level System/36 last week and cut prices on some System/38 models and upgrades. Other System/38 upgrade prices were increased.

In addition, IBM quietly began a new leasing program in which System/36 and 38 customers may break their leases with reduced penalties if they switch to a machine using a different IBM technology. The program is seen as a way of attracting customers to System/36 and 38 processors before the 1988 rollout of the so-called Silverlake machine, which IBM describes as the System/3X follow-on.

The new System/36, the Model 5363, essentially replaces the System/36 PC (Model 5364), although IBM will still offer the older machine.

The 5363 is similar to the 5364 but includes an embedded personal computer and can have the SSP operating system installed at the factory. With the 5364, a user had to attach an IBM Personal Computer or compatible machine or a Personal System/2 to the computer in order to intall the operating system and use the processor.

IBM said performance of the 5363 falls between that of the 5364 and the 5362.

The 5363 is offered with 1M byte of memory standard, compared with the standard amount of 256K bytes with the 5364. However, while the 5364 can grow to 1M byte, the 5363 cannot grow beyond that. The 5363 can handle 28 local displays or printers, compared with a maximum of 16 for the 5364.

"Without the PC, it makes it a lot easier. For users, it will be worth the extra money to have the operating system included," said John McGilvreay, director of customer services at K & C Systems, Inc., a Woburn, Mass.-based development firm specializing in the System/36 PC.

The 5363 can contain one or two 65M-byte disk drives or one or two 105M-byte drives for a maximum of 210M bytes, which exceeds the 5364 maximum disk storage of 130M bytes. The 5363 disk drives come mounted in the system cabinet. A user cannot connect external direct-access storage devices to the processor, IBM said.

Available this Friday, the 5363 starts at $10,000 without the preloaded operating system and runs to about $15,000 for a typically configured system, IBM said.

System/38 price shuffle

In reshuffling prices for the System/38, IBM announced decreases for Models 200 and 700 ranging from 4% to 20%, as well as decreases for memory and upgrades. The price changes became effective last Tuesday.

IBM cut System/38 memory prices from $5,000 to $3,500 per megabyte. IBM also raised the prices for System/38 upgrades from the original family of System/38 Models 5, 7 and 8 to enhanced Models 200, 300 and 400. The upgrade increases range from about 2% on some models to more than 20% on others.

IBM cut prices on upgrades to System/38 Models 600 and 700 up to 20%. Upgrades to a Model 600 must be from a Model 6, 8, 20 or 400. Upgrades to a Model 700 must be from a Model 6, 8, 20, 40, 400 or 600.

The new leasing plan, called the Technology Exchange Option, began last Tuesday. The option is available to all System/38 customers and users of the System/36 5360 and 5362 models. Under a three-year lease from IBM, users of those systems can exchange their processor for a different IBM technology after one year, an IBM spokesman said.

The Model 5363

A compact entry-level System/36 that IBM said can be installed and up and running within one hour

Still offered 5364	New 5363
Separate PC required to load operating system	Preloaded operating system PC components embedded in product Can use with System/36 display, PC or PS/2
256K bytes to 1M byte	1M byte of main storage
16 displays or printers	Up to 28 displays or printers
Up to 130M-byte disk	Up to 210M-byte disk
Price (excluding required PC and operating system): $5,995 to $11,195	Prices (including preloaded operating system): $10,995 (65M-byte disk); $12,095 (105M-byte disk)
Three-month warranty Mid-Range System Amendment not available (no maintenance discount)	One-year warranty Mid-Range System Amendment available (maintenance discount)

CW CHART

Gaps filled in

FROM PAGE 1

At the University of Quebec, data processing director Jean Lapointe said the availability of CICS under VM will allow the school to move some of its existing CICS applications onto its 9370 processor.

IBM also announced VM/SP Release 6, the first VM operating system with the capability to share IBM Conversational Monitor System (CMS) files both within a VM operating system and across VM systems.

Theoretically, VM/SP 6 on a mainframe opens the possibility of CMS file sharing with a 9370 running VM/IS, the version of the operating system designed for the machine, Tasker said. VM/SP 6 is due in December 1988.

'Thorn in the side'

"The lack of file sharing has been a long-standing thorn in the side of VM users," said Gabe Goldberg, director of technology at the VM Systems Group, Inc. in Arlington, Va.

Another major piece of the 9370 puzzle filled was VM/IS Release 5.1, with increased ease of installation and remote control features.

Like the first release of VM/IS, which was based on VM/SP 4, the new system's core is made up of VM/SP 5, but it trails the latest mainframe version. It has received added communications and connectivity functions and was designed to work with IBM's enhanced networking product, Netview.

VM/IS 5.1 is required for CICS/VM but is not scheduled to be available before May 1989.

"No one is going to do anything before mid-1989. It's a little bit depressing," said Peter A. Levine, vice-president of the Gartner Group, Inc. in Stamford, Conn.

Promises aren't enough

"IBM says its customers like to do a lot of advance planning. Customers also like to have products when they need them," Levine added.

Until Tuesday's announcement, 9370 users faced the prospect of running CICS under VSE on the new machine in order to do transaction processing. But VSE lacks VM's communications features, so users were likely to run VSE as a guest under VM, a practice that consumes 22% of the machine's resources, Tasker said.

"Nothing they've announced has an overwhelming impact on my business right now," said Sheldon Danto, a business manager for GTE Corp.'s Data Services division, a 9370 user.

Danto said, however, that he liked the announcement's overall direction toward distributed processing.

Have library, will travel

In addition, IBM has included the Callable Services Library in VM/SP 6. The library will replace the assembler subroutines formerly written by programmers to access VM functions and services.

When different sites write their own assembler routines, applications are not portable among sites, Goldberg said.

VM/SP 6 also features bimodal CMS programming interfaces that allow a programmer to write applications that can run under both the VM/SP and VM/XA environment.

Goldberg called the feature a "bridge" between the SP and XA worlds that would ease migration to XA.

CICS/VM reportedly will be priced at $1,500 a month or carry a one-time charge of $18,000 to $72,000, depending on processor size.

VM/SP 6 will be priced at $500 a month or carry a one-time charge of $7,740 to $30,950, depending on processor size.

VM/IS 5.1 will be priced at $2,381 a month or carry a one-time charge of $28,200 to $106,620.

In addition, IBM announced Version 2 Release 1.7 and Version 3 Release 1.2 of VSE, providing support for IBM 4381 and 3380 processors and peripherals.

Both are due in March.

Version 2 is to carry a monthly license fee of $2,608 or a one-time charge of $29,315 to $102,740.

Version 3 will cost $2,608 per month or $23,110 to $76,475 for a one-time purchase.

Both are priced the same as previous versions.

IBM also announced that it is trying to accelerate the availability of its IMS/VS Version 2.2 Extended Recovery Facility, which provides improved recovery capabilities for large IMS users, by replacing the Early Support Program for the product with managed availability.

Managed availability means selected sites will start receiving the product Nov. 16.

No announcement on general availability will be made in the first quarter of 1988, as previously planned, IBM said.

Netview

FROM PAGE 1

"It significantly expands the number of places where using 9370 is the right solution," said Bob Anderson, IBM's marketing manager for Netview.

Netview/PC Version 1.1, announced in June, was enhanced to support PS/2 Models 50, 60 and 80. Enhancements allow the machines to serve as Netview nodes, reporting data that has been gathered on attached network devices to a Netview host.

IBM also announced the Netview Distribution Manager (DM), which allows a central network operator to distribute data and software changes in an SNA network under IBM's MVS/XA, MVS/370 or VM/SP. The product supersedes the company's Distributed Systems Executive (DSX).

Netview DM can support a VM host, whereas DSX could not. With the new program, a 9370 running VM can use DM to distribute software to other systems. Netview DM also supports the following network extensions: VM end nodes, IBM System/36s acting as intermediate nodes, System/36 end nodes connected to a central site via System/36 intermediate nodes and IBM PC-DOS end nodes connected to a central site via System/36 intermediate nodes.

Netview DM is slated to be available for MVS hosts in June 1988, for VM nodes with MVS hosts in September 1988 and for VM hosts in November 1988.

Analyst lets monitors talk

IBM also announced the Netview Performance Monitor (NPM) Version 1 Release 3, a Netview version of the firm's NPM, which monitors, collects and analyzes network performance information. Due in December, it adds VM support and allows two performance monitors to talk to one another.

"This has been around a long time, but it was not part of Netview. Customers did not have access to it through their Netview console," said Frank Daubeck, president of Communications Network Architects, Inc.

IBM announced the 5822 Data Service Unit/Channel Service Unit (DSU/CSU) for digital networks and a Netview interface for that product. DSU/CSU should be available in second-quarter 1988. IBM also introduced Netview support for its 3737 Channel-to-Channel unit.

The enhancements are priced separately from other Netview components and are not offered as upgrades to current Netview users, IBM said. Pricing is graduated according to the size of the processor.

"These are necessary missing entities. These tweakings of Netview are going to [illegible]

Page from the October 26th, 1987 edition of COMPUTERWORLD:

- "IBM also announced VM/SP Release 6, the first VM operating system with the capability to share IBM … CMS files both within a VM operating system and across VM systems."
- "The lack of file sharing has been a long-standing thorn in the side of VM users," said Gabe Goldberg, director of technology at the VM Systems Group, Inc. in Arlington, VA.

State of the Industry, second half of the 1980's:

- By the way…
  - "Trade press" meant "ink on paper, delivered by the mail carrier."
  - …and there were lots of rumbling noises about this "Unix" thing going on in the background:

State of the Industry, second half of the 1980's:

Andrew Tannenbaum, "Political History of UNIX" (25 Feb '84)

THE CABBAGE PATCH OPERATING SYSTEM

"I sometimes think of UNIX as the Cabbage Patch Operating System.

UNIX wasn't a completely new idea, it was an amalgam of good ideas.

It's certainly the latest craze, with people foaming at the mouth lining up in crazed hordes to get a peek, sometimes paying ridiculous prices for an opportunity to use the product. Of course, they say, it's worth it, there are lesser pleasures in life which are far more expensive.

Like the Cabbage Patch Kids, UNIX has existed for a while, lying around in a relatively dormant state, waiting for the market to explode." …

State of the Industry, second half of the 1980's:

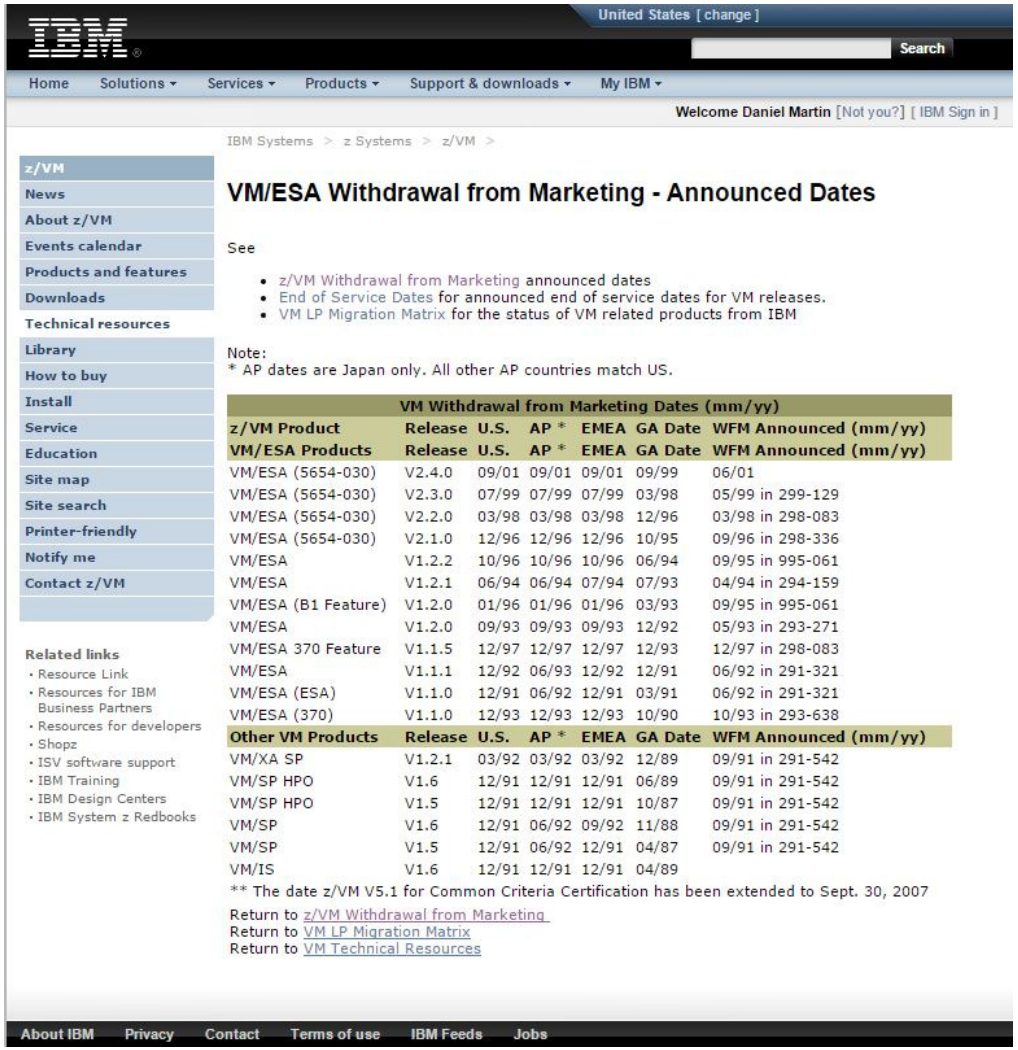

United States [change]

Search

Home | Solutions | Services | Products | Support & downloads | My IBM

Welcome Daniel Martin [Not you?] [IBM Sign in]

IBM Systems > z Systems > z/VM >

z/VM
News
About z/VM
Events calendar
Products and features
Downloads
Technical resources
Library
How to buy
Install
Service
Education
Site map
Site search
Printer-friendly
Notify me
Contact z/VM

Related links
- Resource Link
- Resources for IBM Business Partners
- Resources for developers
- Shopz
- ISV software support
- IBM Training
- IBM Design Centers
- IBM System z Redbooks

VM/ESA Withdrawal from Marketing - Announced Dates

See

- z/VM Withdrawal from Marketing announced dates
- End of Service Dates for announced end of service dates for VM releases.
- VM LP Migration Matrix for the status of VM related products from IBM

Note:
* AP dates are Japan only. All other AP countries match US.

VM Withdrawal from Marketing Dates (mm/yy)						
z/VM Product	Release	U.S.	AP *	EMEA	GA Date	WFM Announced (mm/yy)
VM/ESA Products	Release	U.S.	AP *	EMEA	GA Date	WFM Announced (mm/yy)
VM/ESA (5654-030)	V2.4.0	09/01	09/01	09/01	09/99	06/01
VM/ESA (5654-030)	V2.3.0	07/99	07/99	07/99	03/98	05/99 in 299-129
VM/ESA (5654-030)	V2.2.0	03/98	03/98	03/98	12/96	03/98 in 298-083
VM/ESA (5654-030)	V2.1.0	12/96	12/96	12/96	10/95	09/96 in 298-336
VM/ESA	V1.2.2	10/96	10/96	10/96	06/94	09/95 in 995-061
VM/ESA	V1.2.1	06/94	06/94	07/94	07/93	04/94 in 294-159
VM/ESA (B1 Feature)	V1.2.0	01/96	01/96	01/96	03/93	09/95 in 995-061
VM/ESA	V1.2.0	09/93	09/93	09/93	12/92	05/93 in 293-271
VM/ESA 370 Feature	V1.1.5	12/97	12/97	12/97	12/93	12/97 in 298-083
VM/ESA	V1.1.1	12/92	06/93	12/92	12/91	06/92 in 291-321
VM/ESA (ESA)	V1.1.0	12/91	06/92	12/91	03/91	06/92 in 291-321
VM/ESA (370)	V1.1.0	12/93	12/93	12/93	10/90	10/93 in 293-638
Other VM Products	Release	U.S.	AP *	EMEA	GA Date	WFM Announced (mm/yy)
VM/XA SP	V1.2.1	03/92	03/92	03/92	12/89	09/91 in 291-542
VM/SP HPO	V1.6	12/91	12/91	12/91	06/89	09/91 in 291-542
VM/SP HPO	V1.5	12/91	12/91	12/91	10/87	09/91 in 291-542
VM/SP	V1.6	12/91	06/92	09/92	11/88	09/91 in 291-542
VM/SP	V1.5	12/91	06/92	12/91	04/87	09/91 in 291-542
VM/IS	V1.6	12/91	12/91	12/91	04/89	

** The date z/VM V5.1 for Common Criteria Certification has been extended to Sept. 30, 2007

Return to z/VM Withdrawal from Marketing
Return to VM LP Migration Matrix
Return to VM Technical Resources

About IBM | Privacy | Contact | Terms of use | IBM Feeds | Jobs

Busy times in VM Development:

- VM/SP 1.5 – GA 04/87
- VM/SP 1.6 – GA 11/88
- VM/IS 1.6 – GA 04/89
- VM/XA SP 1.2.1 – GA 12/89

State of the Industry, second half of the 1980's:

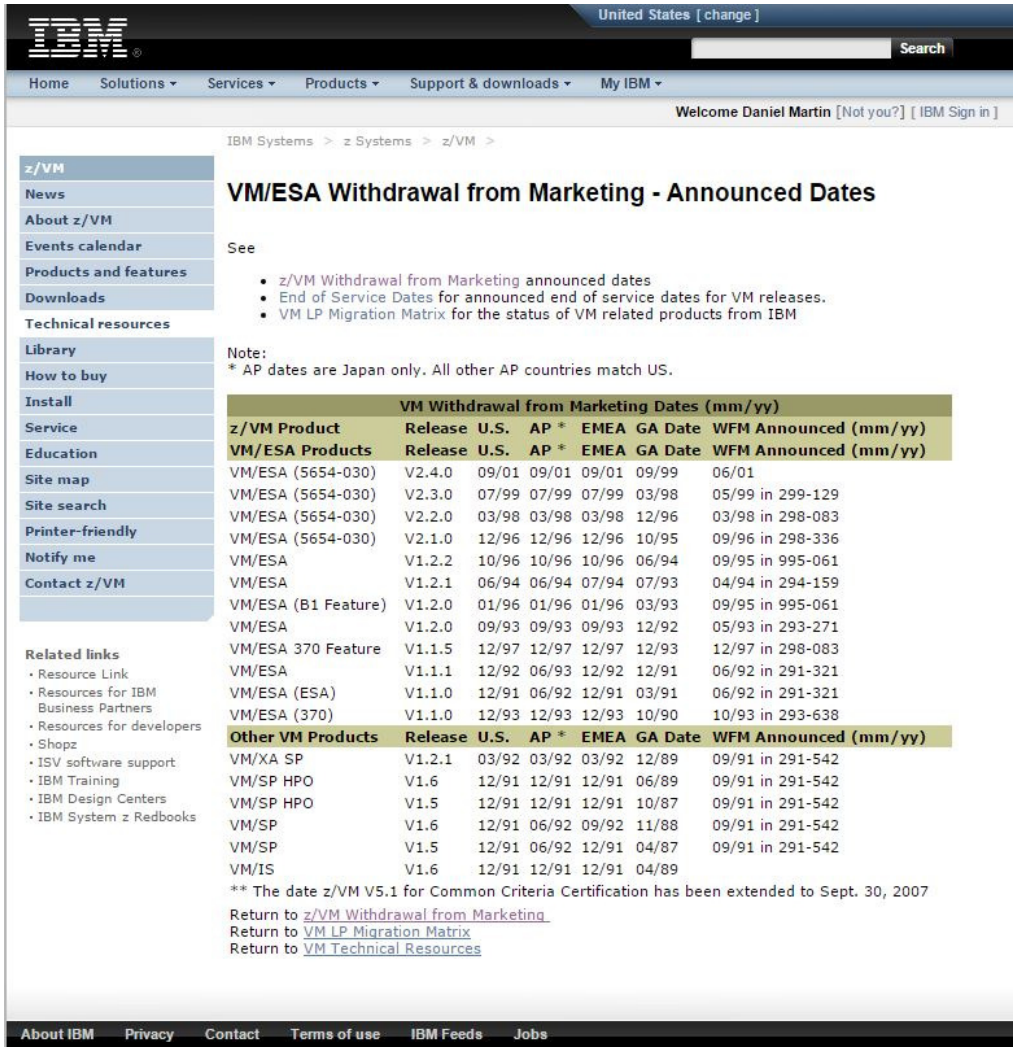

United States [change]

Search

Home | Solutions | Services | Products | Support & downloads | My IBM

Welcome Daniel Martin [Not you?] [IBM Sign in]

IBM Systems > z Systems > z/VM >

z/VM
News
About z/VM
Events calendar
Products and features
Downloads
Technical resources
Library
How to buy
Install
Service
Education
Site map
Site search
Printer-friendly
Notify me
Contact z/VM

Related links
- Resource Link
- Resources for IBM Business Partners
- Resources for developers
- Shopz
- ISV software support
- IBM Training
- IBM Design Centers
- IBM System z Redbooks

VM/ESA Withdrawal from Marketing - Announced Dates

See

- z/VM Withdrawal from Marketing announced dates
- End of Service Dates for announced end of service dates for VM releases.
- VM LP Migration Matrix for the status of VM related products from IBM

Note:
* AP dates are Japan only. All other AP countries match US.

VM Withdrawal from Marketing Dates (mm/yy)						
z/VM Product	Release	U.S.	AP *	EMEA	GA Date	WFM Announced (mm/yy)
VM/ESA Products	Release	U.S.	AP *	EMEA	GA Date	WFM Announced (mm/yy)
VM/ESA (5654-030)	V2.4.0	09/01	09/01	09/01	09/99	06/01
VM/ESA (5654-030)	V2.3.0	07/99	07/99	07/99	03/98	05/99 in 299-129
VM/ESA (5654-030)	V2.2.0	03/98	03/98	03/98	12/96	03/98 in 298-083
VM/ESA (5654-030)	V2.1.0	12/96	12/96	12/96	10/95	09/96 in 298-336
VM/ESA	V1.2.2	10/96	10/96	10/96	06/94	09/95 in 995-061
VM/ESA	V1.2.1	06/94	06/94	07/94	07/93	04/94 in 294-159
VM/ESA (B1 Feature)	V1.2.0	01/96	01/96	01/96	03/93	09/95 in 995-061
VM/ESA	V1.2.0	09/93	09/93	09/93	12/92	05/93 in 293-271
VM/ESA 370 Feature	V1.1.5	12/97	12/97	12/97	12/93	12/97 in 298-083
VM/ESA	V1.1.1	12/92	06/93	12/92	12/91	06/92 in 291-321
VM/ESA (ESA)	V1.1.0	12/91	06/92	12/91	03/91	06/92 in 291-321
VM/ESA (370)	V1.1.0	12/93	12/93	12/93	10/90	10/93 in 293-638
Other VM Products	Release	U.S.	AP *	EMEA	GA Date	WFM Announced (mm/yy)
VM/XA SP	V1.2.1	03/92	03/92	03/92	12/89	09/91 in 291-542
VM/SP HPO	V1.6	12/91	12/91	12/91	06/89	09/91 in 291-542
VM/SP HPO	V1.5	12/91	12/91	12/91	10/87	09/91 in 291-542
VM/SP	V1.6	12/91	06/92	09/92	11/88	09/91 in 291-542
VM/SP	V1.5	12/91	06/92	12/91	04/87	09/91 in 291-542
VM/IS	V1.6	12/91	12/91	12/91	04/89	

** The date z/VM V5.1 for Common Criteria Certification has been extended to Sept. 30, 2007

Return to z/VM Withdrawal from Marketing
Return to VM LP Migration Matrix
Return to VM Technical Resources

About IBM | Privacy | Contact | Terms of use | IBM Feeds | Jobs

Busy times in VM Development:

- VM/SP 1.5 – GA 04/87
- VM/SP 1.6 – GA 11/88
- VM/IS 1.6 – GA 04/89
- VM/XA SP 1.2.1 – GA 12/89

That's FOUR operating system releases, three different packagings, and two radically different hardware architectures shipped inside 32 months.

I was a Higher Ed. customer in the late 1980's. Perhaps I mentioned "stability issues"…

About those stability issues:
VM Development fixed them. A long time ago.

Meanwhile, on the horizon of the early 1990's…

From: torvalds@klaava.Helsinki.FI (Linus Benedict Torvalds)
Newsgroups: comp.os.minix
Subject: What would you like to see most in minix?
Summary: small poll for my new operating system
Message-ID: <1991Aug25.205708.9541@klaava.Helsinki.FI>
Date: 25 Aug 91 20:57:08 GMT
Organization: University of Helsinki

Hello everybody out there using minix –

I'm doing a (free) operating system (just a hobby, won't be big and professional like gnu) for 386(486) AT clones. This has been brewing since april, and is starting to get ready. I'd like any feedback on things people like/dislike in minix, as my OS resembles it somewhat (same physical layout of the file-system (due to practical reasons) among other things).

I've currently ported bash(1.08) and gcc(1.40), and things seem to work. This implies that I'll get something practical within a few months, and I'd like to know what features most people would want. Any suggestions are welcome, but I won't promise I'll implement them 🙂

Linus (torvalds@kruuna.helsinki.fi)

PS. Yes – it's free of any minix code, and it has a multi-threaded fs. It is NOT protable (uses 386 task switching etc), and it probably never will support anything other than AT-harddisks, as that's all I have :-(.

Linus Benedict Torvalds
25 August, 1991
Newsgroups: comp.os.minix

Subject: small poll for my new operating system

Tannenbaum, re-edited:
"Like the Cabbage Patch Kids, ~~UNIX has existed for a while, lying around in a relatively dormant state~~, **[here comes Linux]** waiting for the market to explode." …

SFS early stability issues: Three key points.

- There was an ***unprecedented*** amount of activity in VM development during the second half of the 1980's.
- Linux didn't appear in any form until 1991.
- It's been almost 30 years. Get over it.

Why SFS?

- Consider the times:
  - VM was the platform of choice for “Enterprise Desktop Computing”:
    - …starting with PROFS (PRofessional OFfice System) in 1981 (or earlier, via PRPQ…)
    - …succeeded by OfficeVision/VM
      - (Which wasn’t withdrawn from service until 30 Sept 2003)

Why SFS?

- DASD subsystems were (and are) relatively expensive…
- One CMS virtual machine required a minimum commitment of **one** cylinder of DASD
  - Whether it was ever used or not…
- Multiply by hundreds of systems, and thousands of users…
- And remember, VM had no robust file sharing mechanism.

Solution: SFS

- "CMS file system, re-imagined:"
  - Meets file sharing, security and audit requirements already mentioned.
  - Allows "over-commitment" of DASD capacity.
  - Can be distributed among multiple VM systems.
  - Backward, binary compatibility with existing applications.

- Coincidentally…
  - Enables support for OpenExtensions POSIX file systems.

SFS:
It's Already In The House!

- z/VM 6.3 ships with four SFS File Pool servers pre-installed:

SFS:
z/VM 6.3 Pre-Installed File Pool Servers

- VMSYS:
  - SVM: VMSERVS
  - "System Owned"
  - Used by:
    - LDAP
    - SMAPI
    - SSL
    - DFSMS
    - OE/POSIX
      - *root file system*

SFS:
z/VM 6.3 Pre-Installed File Pool Servers

- VMSYSU:
  - SVM: VMSERVU
  - “System Owned”
  - Used by:
    - z/VM Products
    - z/VM Features
    - DFSMS
    - OE/POSIX:
      - /etc
      - /var
      - /tmp

SFS:
z/VM 6.3 Pre-Installed File Pool Servers

- VMSYSR:
  - SVM: VMSERVR
  - "System Owned"
  - Used as:
    - CRR (Coordinated Resource Recovery) Server
    - *(File under "advanced topics")*

SFS:
z/VM 6.3 Pre-Installed File Pool Servers

- VMPSFS:
  - SVM: VMSERVP
  - "Product Service Data" File Pool
  - Used by:
    - DirMaint
    - SMAPI
    - *"other duties as assigned"*

z/VM 6.3 Pre-Installed File Pool Servers Things to Know:

- "SVM" means "Service Virtual Machine Name," "File Pool Server," or just "server"
- "File Pool Name" is the SFS resource provided by the file pool SVM
- If a file pool name begins with the string "VMSYS" it is a **LOCAL** resource; it will not be shared with other z/VM systems.
- The **VMPSFS:** file pool is defined as a **GLOBAL** resource; it may be shared among multiple z/VM LPARs or SSI members.

z/VM 6.3 Pre-Installed File Pool Servers Things to Know:

- This is a very quick ride through a very deep subject. If you are new to this topic, please consider these lab sessions:

- **TODAY:** Monday, 16:30-18:30: Sessions 17471 & 17472, "*z/VM for Beginners Hands-On Lab*" (two parts, in room Asia 5)

- **TOMORROW:** Tuesday, 13:45-17:30: Sessions 17468, 17469, & 17470, "*z/VM Installation / Migration / Upgrade Hands-On Lab*" (three parts, in room Asia 5)

SFS:
Taking charge, and looking under the hood

Johann Wilhelm Cordes - Wilde Jagd (The Wild Hunt), circa 1856-57

SFS:
Taking charge, and looking under the hood

- Personal Bias:
  - Leave the IBM-supplied file pools alone. They're part of the operating system package.
  - It's your system. The "your gun, your foot, your bullet…" rule applies. You have a policy for that, right?
  - System Programmer / Administrator? Then you're responsible for knowing basic startup, shutdown, health check, and backup/recovery procedures. It's in the documentation.

SFS:
Taking charge, and looking under the hood

- By way of example: setting up a new file pool server.

- Why?
  - Site-specific needs.
    - What's your policy?

  - Installation of additional IBM or 3rd-party program products.
  - Example:
    - IBM Backup and Restore Manager for z/VM maintains backup catalog data in an SFS file space.
    - Use, migration, and disaster recovery may be simplified by provisioning a dedicated file pool server.

SFS:
Taking charge, and looking under the hood

- Setting up a new file pool server:

  - **The Book:** "z/VM: CMS File Pool Planning, Administration, and Operation"

  - At present, the z/VM 6.3 library ships with the z/VM 6.2 version of The Book.
    - This is normal; if a manual is not updated as part of a new release, it's not re-titled.
    - Updates happen. Check the z/VM library: http://www.vm.ibm.com/library/zvmpdf.html

SFS:
Taking charge, and looking under the hood

- CMS File Pool Planning, Administration, and Operation – Chapter 15

- Loosely…
  - Estimate numbers for MAXUSERS and MAXDISKS
  - Create a new file pool service virtual machine
  - Allocate resources, create parameter files
  - Generate the File Pool
  - Start the File Pool
  - Enroll users
  - Proceed with use as needed

SFS:
Taking charge, and looking under the hood

- Documentation overload?
  - "There's a lot of stuff in this chapter…"

- Cheat.
  - The CP Directory entry and configuration files for VMSERVP offer a good "starter template" for defining a new, generic file pool server.
  - VMSERVS and VMSERVU are worth a look, but MDISK sizes are generally too small to be useful.

- FIRST RULE OF SYSTEMS PROGRAMMING:
  - ***"Never go anywhere you don't know how to get back from."***

SFS:
Taking charge, and looking under the hood

```
VMSERVP  DIRECT   A0  F 80  Trunc=72 Size=38 Line=0 Col=1 Alt=0
Editing existing file...

===== * * * Top of File * * *
      T...T....T....T....T....T....T....T....T....T....T....T....T
===== USER VMSERVP RACFPW 64M 64M BG
=====    INCLUDE IBMDFLT
=====    ACCOUNT 1 VMSERVP
=====    IPL 190
=====    IUCV ALLOW
=====    IUCV *IDENT RESANY GLOBAL
=====    MACH XC
=====    OPTION MAXCONN 2000 NOMDCFS APPLMON ACCT QUICKDSP SVMSTAT
=====    POSIXOPT SETIDS ALLOW
=====    SHARE REL 1500
=====    XCONFIG ADDRSPACE MAXNUMBER 100 TOTSIZE 8192G SHARE
=====    XCONFIG ACCESSLIST ALSIZE 1022
=====    CONSOLE 0009 3215 T MAINT
=====    LINK MAINT 0193 0193 RR
=====    MDISK 0191 3390 0632 003 630CM1 WR RSERVER WSERVER
=====    MDISK 0302 3390 0635 015 630CM1 WR RLOG1 WLOG1
=====    MINIOPT NOMDC
=====    MDISK 0301 3390 0650 060 630CM1 WR RCONTROL WCONTROL
=====    MINIOPT NOMDC
=====    MDISK 0303 3390 0710 015 630CM1 WR RLOG2 WLOG2
=====    MINIOPT NOMDC
=====    MDISK 0304 3390 0725 160 630CM1 WR RCATALOG WCATALOG
=====    MDISK 0305 3390 0885 400 630CM1 WR RDATA WDATA
=====    MDISK 0306 3390 1285 400 630CM1 WR RDATA WDATA
=====    MDISK 0307 3390 1685 400 630CM1 WR RDATA WDATA
=====    MDISK 0308 3390 2085 400 630CM1 WR RDATA WDATA
=====    MDISK 0309 3390 2485 400 630CM1 WR RDATA WDATA
=====    MDISK 0310 3390 2885 400 630CM1 WR RDATA WDATA
=====    MDISK 0311 3390 0001 400 630CM2 WR RDATA WDATA
```

Yes, this is a truly awful eye exam.

Let's break it down.

SFS:
Taking charge, and looking under the hood

```
VMSERVP  DIRECT   A0  F 80  Trunc=72 Size=38 Line=0 Col=1 Alt=0
Editing existing file...

===== * * * Top of File * * *
      T...T....T....T....T....T....T....
===== USER VMSERVP RACFPW 64M 64M BG
=====    INCLUDE IBMDFLT
=====    ACCOUNT 1 VMSERVP
=====    IPL 190
=====    IUCV ALLOW
=====    IUCV *IDENT RESANY GLOBAL
=====    MACH XC
=====    OPTION MAXCONN 2000 NOMDCFS APPLMON ACCT QUICKDSP SVMSTAT
=====    POSIXOPT SETIDS ALLOW
=====    SHARE REL 1500
=====    XCONFIG ADDRSPACE MAXNUMBER 100 TOTSIZE 8192G SHARE
=====    XCONFIG ACCESSLIST ALSIZE 1022
=====    CONSOLE 0009 3215 T MAINT
```

Default VMSERVP service virtual machine attributes.

SFS:
Taking charge, and looking under the hood

```
MYOWNSFS  DIRECT   A0  F 80  Trunc=72 Size=38 Line=0 Col=1 Alt=0
Editing existing file...

===== * * * Top of File * * *
      T...T....T....T....T....T....T....T
===== USER MYOWNSFS RACFPW 64M 64M G
=====    INCLUDE IBMDFLT
=====    ACCOUNT 1 MYOWNSFS
=====    IPL 190
=====    IUCV ALLOW
=====    IUCV *IDENT RESANY GLOBAL
=====    MACH XC
=====    OPTION MAXCONN 2000 NOMDCFS APPLMON ACCT QUICKDSP SVMSTAT
=====    POSIXOPT SETIDS ALLOW
=====    SHARE REL 1500
=====    XCONFIG ADDRSPACE MAXNUMBER 100 TOTSIZE 8192G SHARE
=====    XCONFIG ACCESSLIST ALSIZE 1022
=====    CONSOLE 0009 3215 T MAINT
```

Modified CP Directory entry “MYOWNSFS”

-- Based on VMSERVP

-- NEARLY identical to VMSERVP attributes

-- EXCEPT: CP Privilege Class ‘B’ was used by VMSERVP, but removed here.

-- Refer to “Define a Server Machine,” beginning with “Statement 1:” in The Book.

SFS:
Taking charge, and looking under the hood

```
MYOWNSFS DIRECT   A0  F 80  Trunc=72 Size=38 Line=0 Col=1 Alt=0
Editing existing file...

…
=====    LINK MAINT 0193 0193 RR
=====    MDISK 0191 3390 0632 003 630CM1 WR RSERVER WSERVER
=====    MDISK 0302 3390 0635 015 630CM1 WR RLOG1 WLOG1
=====    MINIOPT NOMDC
=====    MDISK 0301 3390 0650 060 630CM1 WR RCONTROL WCONTROL
=====    MINIOPT NOMDC
=====    MDISK 0303 3390 0710 015 630CM1 WR RLOG2 WLOG2
=====    MINIOPT NOMDC
=====    MDISK 0304 3390 0725 160 630CM1 WR RCATALOG WCATALOG
=====    MDISK 0305 3390 0885 400 630CM1 WR RDATA WDATA
=====    MDISK 0306 3390 1285 400 630CM1 WR RDATA WDATA
=====    MDISK 0307 3390 1685 400 630CM1 WR RDATA WDATA
=====    MDISK 0308 3390 2085 400 630CM1 WR RDATA WDATA
=====    MDISK 0309 3390 2485 400 630CM1 WR RDATA WDATA
=====    MDISK 0310 3390 2885 400 630CM1 WR RDATA WDATA
=====    MDISK 0311 3390 0001 400 630CM2 WR RDATA WDATA
```

Default MDISK definitions for VMSERVP.
DO NOT RE-USE THE SAME EXTENTS;
MDISK OVERLAPS WILL DESTROY DATA.

SFS:
Taking charge, and looking under the hood

```
MYOWNSFS DIRECT   A0   F 80   Trunc=72 Size=38 Line=0 Col=1 Alt=0
Editing existing file...

…
=====    LINK MAINT 0193 0193 RR
=====    MDISK 0191 3390 0632 003 630CM1 WR RSERVER WSERVER
=====    MDISK 0302 3390 0635 015 630CM1 WR RLOG1 WLOG1
=====    MINIOPT NOMDC
=====    MDISK 0301 3390 0650 060 630CM1 WR RCONTROL WCONTROL
=====    MINIOPT NOMDC
=====    MDISK 0303 3390 0710 015 630CM1 WR RLOG2 WLOG2
=====    MINIOPT NOMDC
=====    MDISK 0304 3390 0725 160 630CM1 WR RCATALOG WCATALOG
=====    MDISK 0305 3390 0885 400 630CM1 WR RDATA WDATA
=====    MDISK 0306 3390 1285 400 630CM1 WR RDATA WDATA
=====    MDISK 0307 3390 1685 400 630CM1 WR RDATA WDATA
=====    MDISK 0308 3390 2085 400 630CM1 WR RDATA WDATA
=====    MDISK 0309 3390 2485 400 630CM1 WR RDATA WDATA
=====    MDISK 0310 3390 2885 400 630CM1 WR RDATA WDATA
=====    MDISK 0311 3390 0001 400 630CM2 WR RDATA WDATA
```

Read-only LINK to MAINT 193 – SFS run-time commands and utilities

SFS:
Taking charge, and looking under the hood

```
MYOWNSFS DIRECT   A0   F 80   Trunc=72 Size=38 Line=0 Col=1 Alt=0
Editing existing file...

…
=====    LINK MAINT 0193 0193 RR
=====    MDISK 0191 3390 nnnn 003 volser WR RSERVER WSERVER
=====    MDISK 0302 3390 nnnn 015 volser WR RLOG1 WLOG1
=====    MINIOPT NOMDC
=====    MDISK 0301 3390 nnnn 060 volser WR RCONTROL WCONTROL
=====    MINIOPT NOMDC
=====    MDISK 0303 3390 nnnn 015 volser WR RLOG2 WLOG2
=====    MINIOPT NOMDC
=====    MDISK 0304 3390 nnnn 160 volser WR RCATALOG WCATALOG
=====    MDISK 0305 3390 nnnn 400 volser WR RDATA WDATA
=====    MDISK 0306 3390 nnnn 400 volser WR RDATA WDATA
=====    MDISK 0307 3390 nnnn 400 volser WR RDATA WDATA
=====    MDISK 0308 3390 nnnn 400 volser WR RDATA WDATA
=====    MDISK 0309 3390 nnnn 400 volser WR RDATA WDATA
=====    MDISK 0310 3390 nnnn 400 volser WR RDATA WDATA
=====    MDISK 0311 3390 nnnn 400 volser WR RDATA WDATA
```

MDISK 191 – File Pool Server configuration files.

SFS:
Taking charge, and looking under the hood

```
MYOWNSFS DIRECT   A0   F 80   Trunc=72 Size=38 Line=0 Col=1 Alt=0
Editing existing file...

…
=====    LINK MAINT 0193 0193 RR
=====    MDISK 0191 3390 nnnn 003 volser WR RSERVER WSERVER
=====    MDISK 0302 3390 nnnn 015 volser WR RLOG1 WLOG1
=====    MINIOPT NOMDC
=====    MDISK 0301 3390 nnnn 060 volser WR RCONTROL WCONTROL
=====    MINIOPT NOMDC
=====    MDISK 0303 3390 nnnn 015 volser WR RLOG2 WLOG2
=====    MINIOPT NOMDC
=====    MDISK 0304 3390 nnnn 160 volser WR RCATALOG WCATALOG
=====    MDISK 0305 3390 nnnn 400 volser WR RDATA WDATA
=====    MDISK 0306 3390 nnnn 400 volser WR RDATA WDATA
=====    MDISK 0307 3390 nnnn 400 volser WR RDATA WDATA
=====    MDISK 0308 3390 nnnn 400 volser WR RDATA WDATA
=====    MDISK 0309 3390 nnnn 400 volser WR RDATA WDATA
=====    MDISK 0310 3390 nnnn 400 volser WR RDATA WDATA
=====    MDISK 0311 3390 nnnn 400 volser WR RDATA WDATA
```

MDISKs 302 and 303 – File Pool Server log data minidisks.

SFS:
Taking charge, and looking under the hood

```
MYOWNSFS DIRECT   A0   F 80   Trunc=72 Size=38 Line=0 Col=1 Alt=0
Editing existing file...

…
=====    LINK MAINT 0193 0193 RR
=====    MDISK 0191 3390 nnnn 003 volser WR RSERVER WSERVER
=====    MDISK 0302 3390 nnnn 015 volser WR RLOG1 WLOG1
=====    MINIOPT NOMDC
=====    MDISK 0301 3390 nnnn 060 volser WR RCONTROL WCONTROL
=====    MINIOPT NOMDC
=====    MDISK 0303 3390 nnnn 015 volser WR RLOG2 WLOG2
=====    MINIOPT NOMDC
=====    MDISK 0304 3390 nnnn 160 volser WR RCATALOG WCATALOG
=====    MDISK 0305 3390 nnnn 400 volser WR RDATA WDATA
=====    MDISK 0306 3390 nnnn 400 volser WR RDATA WDATA
=====    MDISK 0307 3390 nnnn 400 volser WR RDATA WDATA
=====    MDISK 0308 3390 nnnn 400 volser WR RDATA WDATA
=====    MDISK 0309 3390 nnnn 400 volser WR RDATA WDATA
=====    MDISK 0310 3390 nnnn 400 volser WR RDATA WDATA
=====    MDISK 0311 3390 nnnn 400 volser WR RDATA WDATA
```

MDISK 301 – File Pool Server control data minidisk.

SFS:
Taking charge, and looking under the hood

```
MYOWNSFS DIRECT   A0   F 80   Trunc=72 Size=38 Line=0 Col=1 Alt=0
Editing existing file...

…
=====    LINK MAINT 0193 0193 RR
=====    MDISK 0191 3390 nnnn 003 volser WR RSERVER WSERVER
=====    MDISK 0302 3390 nnnn 015 volser WR RLOG1 WLOG1
=====    MINIOPT NOMDC
=====    MDISK 0301 3390 nnnn 060 volser WR RCONTROL WCONTROL
=====    MINIOPT NOMDC
=====    MDISK 0303 3390 nnnn 015 volser WR RLOG2 WLOG2
=====    MINIOPT NOMDC
=====    MDISK 0304 3390 nnnn 160 volser WR RCATALOG WCATALOG
=====    MDISK 0305 3390 nnnn 400 volser WR RDATA WDATA
=====    MDISK 0306 3390 nnnn 400 volser WR RDATA WDATA
=====    MDISK 0307 3390 nnnn 400 volser WR RDATA WDATA
=====    MDISK 0308 3390 nnnn 400 volser WR RDATA WDATA
=====    MDISK 0309 3390 nnnn 400 volser WR RDATA WDATA
=====    MDISK 0310 3390 nnnn 400 volser WR RDATA WDATA
=====    MDISK 0311 3390 nnnn 400 volser WR RDATA WDATA
```

MDISK 304 – File Pool Server catalog data minidisk.

SFS:
Taking charge, and looking under the hood

```
MYOWNSFS DIRECT   A0   F 80   Trunc=72 Size=38 Line=0 Col=1 Alt=0
Editing existing file...

…
=====    LINK MAINT 0193 0193 RR
=====    MDISK 0191 3390 nnnn 003 volser WR RSERVER WSERVER
=====    MDISK 0302 3390 nnnn 015 volser WR RLOG1 WLOG1
=====    MINIOPT NOMDC
=====    MDISK 0301 3390 nnnn 060 volser WR RCONTROL WCONTROL
=====    MINIOPT NOMDC
=====    MDISK 0303 3390 nnnn 015 volser WR RLOG2 WLOG2
=====    MINIOPT NOMDC
=====    MDISK 0304 3390 nnnn 160 volser WR RCATALOG WCATALOG
=====    MDISK 0305 3390 nnnn 400 volser WR RDATA WDATA
=====    MDISK 0306 3390 nnnn 400 volser WR RDATA WDATA
=====    MDISK 0307 3390 nnnn 400 volser WR RDATA WDATA
=====    MDISK 0308 3390 nnnn 400 volser WR RDATA WDATA
=====    MDISK 0309 3390 nnnn 400 volser WR RDATA WDATA
=====    MDISK 0310 3390 nnnn 400 volser WR RDATA WDATA
=====    MDISK 0311 3390 nnnn 400 volser WR RDATA WDATA
```

MDISKs 305-311 – File Pool Server end-user data minidisks

MDISKs 305-311 – Size and number of data MDISKs is a local choice. Each 3390 cylinder yields 180 4K file pool data blocks.

SFS:
Taking charge, and looking under the hood

- Continuing along in **The Book**, Chapter 15…

- Estimates for MAXUSERS and MAXDISKS done.

- MYOWNSFS service virtual machine added to CP Directory.

- Log on to MYOWNSFS
- FORMAT the A-Disk
- Create POOLDEF, DMSPARMS, and PROFILE EXEC
- ACCESS the 193 minidisk
- Execute FILESERV GENERATE

SFS:
Taking charge, and looking under the hood

- **Create the Startup Parameter file, *myownsfs DMSPARMS***

```
 MYOWNSFS DMSPARMS A1  V 80  Trunc=80 Size=13 Line=0 Col=1 Alt=0
Editing existing file...
Upper case set.
```

The Book, Chapter 15, Step 7 describes DMSPARMS settings.

```
===== * * * Top of File * * *
      T...T....T....T....T....T....T....T....T....T....T....T....T....T....T....T....T....T....T....T
===== ********************************
===== * MYOWNSFS DMSPARMS (MYOWNSFS) *
===== ********************************
=====
===== ADMIN MAINT MAINT630 AUTOLOG1 AUTOLOG2 MIGMAINT VSMGUARD VSMWORK1
===== ADMIN VSMWORK2 VSMWORK3
===== ADMIN BKRBKUP BKRWRK01 BKRWRK02 BKRWRK03 BKRWRK04 BKRWRK05 BKRWRK06
===== NOBACKUP
===== SAVESEGID CMSFILES
===== REMOTE
===== USERS 100
===== * * * End of File * * *
```

SFS:
Taking charge, and looking under the hood

- Execute **FILESERV GENERATE**

```
 $$TEMP  $POOLDEF  Z1  F 80  Trunc=80 Size=18 Line=0 Col=1 Alt
Editing existing file...
Upper case set.

===== * * * Top of File * * *
      T...T....T....T....T....T....T....T....T....T....T....T.
===== MAXUSERS=1000
===== MAXDISKS=500
===== DDNAME=CONTROL         VDEV=301
===== DDNAME=LOG1            VDEV=302
===== DDNAME=LOG2            VDEV=303
===== DDNAME=BACKUP   DISK   FN=FILEPOOL   FT=BACKUP    FM=*
===== DDNAME=MDK00001        VDEV=304      GROUP=1   BLOCKS=0
===== DDNAME=MDK00002        VDEV=305      GROUP=2   BLOCKS=0
===== DDNAME=MDK00003        VDEV=306      GROUP=2   BLOCKS=0
===== DDNAME=MDK00004        VDEV=307      GROUP=2   BLOCKS=0
===== DDNAME=MDK00005        VDEV=308      GROUP=2   BLOCKS=0
===== DDNAME=MDK00006        VDEV=309      GROUP=2   BLOCKS=0
===== DDNAME=MDK00007        VDEV=310      GROUP=2   BLOCKS=0
===== DDNAME=MDK00008        VDEV=311      GROUP=2   BLOCKS=0
===== * * * End of File * * *
```

The Book, Chapter 15, Step 8 describes POOLDEF settings.

THIS STEP FORMATS MDISKs.

DOUBLE-CHECK YOUR WORK BEFORE INVOKING *FILESERV GENERATE*

SFS:
Care and Feeding of a File Pool Server

- Minimum configuration:
  - Sample **PROFILE EXEC** as shipped with the system
  - PLEASE be sure to **CP SPOOL CONSOLE…** in order to capture the console log.
  - "But there's never anything in the console log."
    - "That's right. Except when something breaks."

SFS:
Care and Feeding of a File Pool Server

- Minimum administrative skill set:
  - Startup
  - Shutdown
  - Basic Administration
  - Backup / Recovery

SFS:
Care and Feeding of a File Pool Server

- Startup
  - Usually via AUTOLOG1 (non-RACF) or AUTOLOG2 (RACF)
  - …or site choice of automated operations tool set.

- Shutdown
  - File Pool SVM Console "**STOP**" command
  - Via **CP SIGNAL SHUTDOWN**
    - …if enabled in DMSPARMS file

SFS Administration: Care and Feeding of a File Pool Server

- Basic Administration
  - ENROLL USER – create an SFS file space
    - Convention: Virtual Machine Name == File Space Name
  - MODIFY USER – alter file space quota limits
  - DELETE USER – destroy an SFS file space

  - ENROLL ADMINISTRATOR
    - Transient; define permanent ADMINs in DMSPARMS
  - DELETE ADMINISTRATOR
    - Transient; permanent ADMINs in DMSPARMS will return when the file pool is restarted.

SFS Administration: Care and Feeding of a File Pool Server

- Backup / Recovery
  - "It's complicated."
  - This is at least as much a policy decision as it is a technology decision.
    - What's the business value of your data today?

- "Big Shop" solutions tend to rely on remote DASD replication. This is fine for "Big Shop" disaster recovery problems.

- …But I only deleted one file…?
  - Well… then it's complicated.

SFS Administration: Care and Feeding of a File Pool Server

- Backup / Recovery

- IBM-supplied tool set:
  - FILEPOOL UNLOAD and FILEPOOL RELOAD
    - Per-file space or Per-storage group
    - These work. The price is right. They're hard to automate.

  - FILEPOOL BACKUP and FILEPOOL RESTORE
    - Per-storage group
    - Bigger tools. Same problems. No "just one file" restore.

SFS Administration: Care and Feeding of a File Pool Server

- Backup / Recovery

- z/VM Native IBM and OEM solutions
  - IBM Backup and Restore Manager for z/VM
  - CA VM:Manager Suite
  - Home-grown?
  - Other?
    - I'm curious… what do YOU do?

SFS Administration: Care and Feeding of a File Pool Server

- Backup / Recovery

- Remote DASD replication
  - Is it worth breaking your hot site link to recover one file?
  - What's your policy for that?
  - What's your replication lag time?
    - "Oops."

SFS Administration:
Care and Feeding of a File Pool Server

- Backup / Recovery

- DASD image backup / off-platform solutions
  - DDR
  - DFDSS, FDR

- These are fine solutions, IF…
  - …You can shut down the file pool server before backing up all DASD assets.
  - ...and IF applications involved don't have "data in flight" exposures.

- z/VM SFS recovery is similar to Linux LVM recovery: image backup is great, **if you can maintain point-in-time consistency across all DASD assets**.

- Did I mention "Policy"?

SFS Usage:
End User Interaction

- "Classic" CMS commands are generally unaware of SFS, for ACCESSed directories
  - XEDIT has built-in awareness; it performs LOCK / UNLOCK operations when appropriate.

- "Old code" backward compatibility is excellent.
  - **ACCESS *vdev* FM**
  - **ACCESS** *filepool:filespace.path1.path2…* **FM**

- Most native CMS commands can refer to either
  - ***CMD filename filetype fm***

    -or-
  - ***CMD filename filetype filepool:filespace.path1.path2...***

SFS Usage:
End User Interaction

- Additional commands
  - CREATE DIRECTORY
    - Two types of SFS directory
      - *FILECONTROL*
        - » File-by-file permission management
        - » File-level sharing, lock management, transaction integrity
      - *DIRCONTROL*
        - » “Works more like a minidisk”
        - » One writer, multiple readers
        - » No file-level permissions
  - Destroyed by ERASE

SFS Usage:
End User Interaction

- Additional commands

- GRANT AUTHORITY / REVOKE AUTHORITY

  - FILECONTROL: READ, WRITE, NEWREAD, NEWWRITE
  - Files in FILECONTROL directories: READ, WRITE

  - DIRCONTROL: READ, WRITE
  - Files in DIRCONTROL directories: determined by directory permissions; no per-file authorizations in DIRCONTROL.

  - "Just different enough…" to surprise the uninitiated:
    - No permission inheritance

SFS Usage:
End User Interaction

- A bit more about permission management:

- Native behavior is subject to External Security Management override if RACF or other ESM is enabled.

- Possible Administration / Audit headache:
  - The native command set makes it difficult to deconstruct complex permission hierarchies.
  - It's difficult to simply say "Give user Wally the same permissions to everything in this file pool as user Alice."
    - Solutions: ESM control, OEM tools, or roll-your-own.
    - WIBNI: "Wouldn't It Be Nice If…"

SFS Usage:
End User Interaction

- File system navigation

- "Old school"
  - LISTFILE and FILELIST operate seamlessly on ACCESSed directories.

- SFS-specific
  - The CMS DIRLIST command provides directory and file-level navigation.
  - …plus permission inquiry.
  - …plus drill-down into FILELIST.

SFS Usage:
End User Interaction – QUERY ACCESSED

- File system navigation:

```
query accessed
Mode  Stat      Files  Vdev  Label/Directory
A      R/W        455  292   DM292Z
B      R/W       2013  DIR   MYOWNSFS:DMARTIN.MAINT190COPY
E/E    R/O        194  191   DM191Z
M      R/O          4  198   TCM198
N      R/O        859  592   TCM592
O      R/O        142  1B5   SMS1B5
S      R/O        697  190   MNT190
W/W    R/O       1311  993   MNT193
X/S    R/O       1222  19D   MNT19D
Y/S    R/O       1554  19E   MNT19E
Ready;
```

```
RUNNING    SHARE125
```

SFS Usage:
End User Interaction - DIRLIST

- File system navigation:

```
 DMARTIN  DIRLIST  A0  V 319  Trunc=319 Size=1 Line=1 Col=1 Alt=0
Cmd   Fm Directory Name
      -  MYOWNSFS:DMARTIN.
      B  MYOWNSFS:DMARTIN.MAINT190COPY
```

```
1= Help       2= Refresh  3= Quit   4= Sort(fm)    5= Sort(dir)   6= Auth
7= Backward   8= Forward  9=       10=            11= Filelist   12= Cursor

====>
```

```
X E D I T  1 File
```

SFS Usage:
End User Interaction – DIRLIST drill-down

- File system navigation:

```
 DMARTIN  FILELIST A0  V 169  Trunc=169 Size=2013 Line=1 Col=1 Alt=0
Directory = MYOWNSFS:DMARTIN.
Cmd   Filename Filetype Fm Format Lrecl     Records     Blocks   Date      Time
      ASSEMBLE MODULE   B2 V        8192           3          3  3/31/15 17:01:57
      IFOX00   MODULE   B2 V        2904           3          1  3/31/15 17:01:57
      IFOX01   MODULE   B2 V        1688           3          1  3/31/15 17:01:57
      IFOX02   MODULE   B2 V        3248           3          1  3/31/15 17:01:57
      IFOX03   MODULE   B2 V         320           3          1  3/31/15 17:01:57
      IFOX04   MODULE   B2 V         832           3          1  3/31/15 17:01:57
      IFOX05   MODULE   B2 V         536           3          1  3/31/15 17:01:57
      IFOX06   MODULE   B2 V        1032           3          1  3/31/15 17:01:57
      IFOX07   MODULE   B2 V        1808           3          1  3/31/15 17:01:57
      IFOX11   MODULE   B2 V       19008           3          5  3/31/15 17:01:57
      IFOX21   MODULE   B2 V        5504           3          2  3/31/15 17:01:57
      IFOX31   MODULE   B2 V       14128           3          4  3/31/15 17:01:57
      IFOX41   MODULE   B2 V        9088           3          3  3/31/15 17:01:57
      IFOX42   MODULE   B2 V        9992           3          3  3/31/15 17:01:57
      IFOX51   MODULE   B2 V       21824           3          6  3/31/15 17:01:57
      IFOX61   MODULE   B2 V        4696           3          2  3/31/15 17:01:57
1= Help       2= Refresh  3= Quit   4= Cancel      5= Sort(dir)   6= Sort(size)
7= Backward   8= Forward  9= FL /n 10= Share      11= XEDIT/LIST 12= Cursor

====>

                                                    X E D I T  1 File
```

SFS Usage:
End User Interaction – FILELIST * * S comparison

- File system navigation:

```
 DMARTIN  FILELIST A0  V 169  Trunc=169 Size=697 Line=1 Col=1 Alt=0
Cmd   Filename Filetype Fm Format Lrecl     Records     Blocks   Date     Time
      ASSEMBLE MODULE   S2 V        8192           3          3  3/31/15 17:01:57
      IFOX00   MODULE   S2 V        2904           3          1  3/31/15 17:01:57
      IFOX01   MODULE   S2 V        1688           3          1  3/31/15 17:01:57
      IFOX02   MODULE   S2 V        3248           3          1  3/31/15 17:01:57
      IFOX03   MODULE   S2 V         320           3          1  3/31/15 17:01:57
      IFOX04   MODULE   S2 V         832           3          1  3/31/15 17:01:57
      IFOX05   MODULE   S2 V         536           3          1  3/31/15 17:01:57
      IFOX06   MODULE   S2 V        1032           3          1  3/31/15 17:01:57
      IFOX07   MODULE   S2 V        1808           3          1  3/31/15 17:01:57
      IFOX11   MODULE   S2 V       19008           3          5  3/31/15 17:01:57
      IFOX21   MODULE   S2 V        5504           3          2  3/31/15 17:01:57
      IFOX31   MODULE   S2 V       14128           3          4  3/31/15 17:01:57
      IFOX41   MODULE   S2 V        9088           3          3  3/31/15 17:01:57
      IFOX42   MODULE   S2 V        9992           3          3  3/31/15 17:01:57
      IFOX51   MODULE   S2 V       21824           3          6  3/31/15 17:01:57
      IFOX61   MODULE   S2 V        4696           3          2  3/31/15 17:01:57
      IFOX62   MODULE   S2 V       15144           3          4  3/31/15 17:01:57
1= Help       2= Refresh  3= Quit   4= Sort(type)  5= Sort(date)  6= Sort(size)
7= Backward   8= Forward  9= FL /n 10=            11= XEDIT/LIST 12= Cursor

====>

                                                    X E DI T  1 File
```

SFS Usage:
End User Interaction - BFS

- Quick peek: OpenExtensions and BFS

- OpenExtensions: POSIX standards under CMS:
  POSIX 1003.1 (POSIX.1) - System Interfaces
  POSIX 1003.1a (POSIX.1a) - Extensions to POSIX.1
  POSIX 1003.1c (POSIX.1c) - Threads
  POSIX 1003.2 (POSIX.2) - Shell and Utilities

- BFS: Byte File System
  - Stream-oriented, POSIX-compliant file system
  - Supported via SFS file pool server
  - “The other, other file system” for CMS

SFS Usage:
End User Interaction

- OpenExtensions / CMS and BFS
- Minimalist environment
  - /bin/sh
  - Bare necessities for POSIX compliance

- Details are documented in:
  - z/VM OpenExtensions POSIX Conformance Document
  - z/VM OpenExtensions User's Guide
  - z/VM OpenExtensions Commands Reference

SFS Usage:
End User Interaction

- OpenExtensions / CMS and BFS

- Traditionally: ported UNIX features
- Today: porting of some z/OS packages to z/VM
  - Because “why develop it twice?”
  - LDAP, SSL, MPROUTE

- From the “Ready;” prompt, invoked by **OPENVM**
  - mount / unmount, create, erase, listfile, set…
    - *See z/VM: OpenExtensions Commands Reference*

SFS Usage: OpenExtensions CMS & BFS

```
Ready;
openvm query mount
Mount point = '/home/dmartin'
Type  Stat  Mounted
BFS   R/W   '/../VMBFS:VMLOCAL:DPMFS/'
Mount point = '/tmp'
Type  Stat  Mounted
BFS   R/W   '/../VMBFS:VMLOCAL:TMPFS/'
Mount point = '/'
Type  Stat  Mounted
BFS   R/W   '/../VMBFS:VMLOCAL:ALTROOT/'
Ready;
```

SFS Usage: OpenExtensions CMS & BFS

```
Ready;
openvm shell

IBM
Licensed Material - Property of IBM
5654-A17 (C) Copyright IBM Corp. 1995
(C) Copyright Mortice Kern Systems, Inc., 1985, 1993.
(C) Copyright Software Development Group, University of Waterloo, 1989.

All Rights Reserved.

U.S. Government users - RESTRICTED RIGHTS - Use, Duplication, or
Disclosure restricted by GSA-ADP schedule contract with IBM Corp.

IBM is a registered trademark of the IBM Corp.


#
pwd
/
#
cd /home/dmartin
```

SFS Usage: OpenExtensions CMS & BFS

```
#
ps -f
     UID        PID       PPID    STIME TTY        TIME COMMAND
 dmartin       3733          1 01:27:43 tty        0:00
 dmartin       4532       3733 01:27:43 tty        0:00 sh -L
 dmartin       4641       4532 01:27:59 tty        0:00 ps -f
#
ls -al
total 16
-rw-r--r--   2 dmartin  dpm           26 Aug 31  2012 hardlinkfile
lrwxrwxrwx   1 dmartin  dpm            7 Aug 31  2012 linkdir -> testdir
lrwxrwxrwx   1 dmartin  dpm            8 Aug 31  2012 linkfile -> testfile
drwxr-xr-x   1 dmartin  dpm            0 May 26 19:43 oddballs
drwxr-xr-x   1 dmartin  dpm            0 Aug 31  2012 testdir
-rw-r--r--   2 dmartin  dpm           26 Aug 31  2012 testfile
#
exit
Ready;
```

SFS Usage: OpenExtensions CMS & BFS

- XEDIT can “cross the boundary” to edit OpenExtensions BFS files
  - From a “Ready;” prompt:
    - xedit /../vmbfs:vmlocal:dpmbfs/testfile (nametype bfs
  - From a shell prompt:
    - cms xedit /../vmbfs:vmlocal:dpmbfs/testfile \(nametype bfs
  - BFS file system **E**xternal constructs allow application-friendly references

Summary:
SFS Usage and Administration

- SFS is:
  - Part of your z/VM System Infrastructure
    - Four servers, right out of the box
  - Centrally involved in your ability to operate and maintain z/VM
    - Product service, SSI, DirMaint, SMAPI, LDAP…
  - A fully viable, fully supported, robust file system for CMS
  - Required by some IBM and OEM products, in addition to native z/VM functions and services

Summary:
SFS Usage and Administration

- Bibliography
  - z/VM: CMS File Pool Planning, Administration, and Operation
  - z/VM: CP Planning and Administration
  - z/VM: CMS Commands and Utilities Reference
  - z/VM: CMS User's Guide
  - z/VM: OpenExtensions POSIX Conformance Document
  - z/VM: OpenExtensions User's Guide
  - z/VM: OpenExtensions Commands Reference

Summary:
SFS Usage and Administration

- Online resources:
  - IBM: www.vm.ibm.com/library/index.html
  - User Community:
    - SHARE! - www.share.org
    - IBMVM Mailing List: http://listserv.uark.edu/archives/ibmvm.html
    - Linux-390 Mailing List: http://www2.marist.edu/htbin/wlvindex?linux-390

SFS Usage and Administration: Questions, comments, and feedback?

THANK YOU! – Session 17274
CMS Shared File System Usage & Administration

Daniel P. Martin
Senior Software Developer
Rocket Software
1400 NE McClain Rd, Suite 8 • Bentonville, AR 72712 • USA
• e: dmartin@rocketsoftware.com
• w: www.rocketsoftware.com

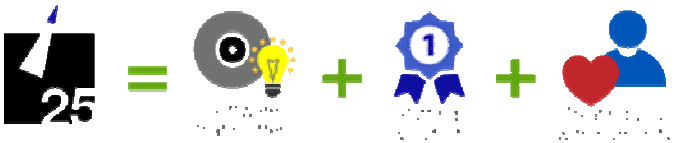

That's not rocket science. That's Rocket Software.

Complete your session evaluations online at www.SHARE.org/Orlando-Eval