

IBM CORPORATION

Moderator: Susan Greenlee
June 15, 2004
1:00 p.m. CT

Operator: Good day, everyone, and welcome to today's IBM z/VM conference. Today's call is being recorded.

At this time, I'd like to turn the conference over to Susan Greenlee. Susan, please go ahead.

Susan Greenlee: Hello, and welcome to the third call in a series of quarterly Linux on zSeries technical education teleconferences, prepared especially for you, our zSeries customers. Thank you for choosing to spend some of your valuable time with us today. My name is Susan Greenlee, and I'm an IBM Linux technical support marketing specialist for zSeries. And it's my pleasure to serve as your moderator for today's session.

Today, we will provide you with an overview of the enhancements to z/VM that were announced on April 7th for z/VM Version 5, Release 1, with planned general availability on September 24th. These enhancements fall into four major categories: first, virtualization technology and Linux enablement; secondly, network virtualization and security; thirdly, technology exploitation; and finally, systems management. In addition, we'll cover, the new engine based value unit pricing for z/VM Version 5, which replaces the per engine pricing model that is available today in z/VM Version 4.

If you've attended VM users group meetings or technical conferences, you may have had the opportunity to hear today's presenter, Romney White from VM Strategy and Development. Romney has worked with VM for more than 30 years. Before joining IBM in 1997, he founded four VM based software companies. He has been instrumental in bringing the potential contribution of VM to a Linux environment and joins us today from the VM development lab in Endicott, New York.

Romney, thank you for joining us on the broadcast today.

Romney White: Thank you, Susan. It's a pleasure to be with you.

Susan Greenlee: I'll turn it over to you Romney in a minute, but first let me cover the flow of the call today, which is the same as what we've done on our previous calls. Romney will present for about 50 minutes. Then, we'll conduct a brief electronic survey. Then, you will have a chance to ask questions about today's material. And at the end, I'll close with directions on how to obtain more information about Linux and a replay of this teleconference.

So, with that, let's get started, Romney.

Romney White: OK, Susan, thanks. Good morning or good afternoon, depending on where you are. It's a pleasure to be here to speak with you today.

I'm going to start on chart two, the agenda, which will cover the evolution of z/VM, talk about terms and conditions, including pricing. But I want to spend most of the time we have today talking about the enhancements in z/VM Version 5, Release 1. We have removed a few functions, and I'll cover those briefly. There are some packaging and delivery changes that I'll also touch on. And we made a few statements of direction that I think it's important you

understand. And finally, I'll talk briefly about futures and things that we're thinking about working on.

On chart three, you'll see a depiction of VM's evolution. You'll notice that z/VM Version 3, which supports some of the older processors in our line, is turning red – that is to say, will be withdrawn from marketing in August of this year. So, the direction for future versions of z/VM is to go to more modern processors. And in fact, Version 5 which we plan to GA in September, as Susan mentioned, will support zSeries processors only. If you need to run on a pre-zSeries processor, a G5 or a G6 9672, then z/VM Release 4 will be available for that purpose.

On chart four, you can see the release frequency showing a comparison between VM/ESA, z/VM's predecessor, and z/VM. And as you'll see, we've increased the frequency of z/VM releases quite dramatically in a few cases. But what we're trying to get to and what our plans are, are to have about one release a year. Now, it won't necessarily be every 12 months, but to the best of our ability, we'll try to come out with a new release about once every year.

Chart five talks about the terms and conditions for Version 5. And again, Susan mentioned we've changed the pricing model. It's what we call value unit based charging. It's still a one-time charge, but the charge is now graduated. And so, essentially, the more you buy, the less expensive it is. We still have subscription support, and we recommend that everybody sign up for this, since this gives you the traditional kind of telephone support that you're used to with the monthly license charge products. It also entitles you to new versions and releases of z/VM. So, if you currently pay subscription and support for z/VM Version 4, for example, you're entitled to z/VM Version 5 at no additional charge.

And IPLA applies not only to z/VM, but also the optional features, which include the directory maintenance product, rack halves for VM, the security product and the new performance toolkit.

Chart six shows the pricing that we're planning to offer for z/VM Version 5 and compares it with what you see with what we offered for z/VM Version 4. The red line, straight line, is for Version 4 pricing. As you can see, it was the same price per engine, no matter how many engines you had. With Version 5, the pricing is per engine, and it's graduated. So, as the chart shows, if you have four to six engines, then you pay \$20,250 per engine. And if you grow into your seventh engine, the z/VM price is \$18,000 and so on. So, the more you have, the less expensive each engine becomes.

As well, you can aggregate engines across your enterprise so that you have multiple systems in your enterprise. You can add up the number of engines you have to arrive at the price per engine. So, we hope this will make it easier for you to acquire z/VM and install it on all of your zSeries processors.

Chart seven lists the enhancement areas for z/VM Version 5, Release 1. And these are the areas in which I'll be concentrating on in this presentation. First, virtualization enhancements – that is, enhancements that make Linux on zSeries run more efficiently and with more function. And these enhancements are listed on chart eight, and we'll go through each of them in turn.

Chart nine introduces the real centerpiece of this release, native support for SCSI disks. The objectives here are not to persuade you to throw out your account key devices and replace them with SCSI volumes, but rather to let you introduce z/VM into an existing SCSI environment.

So, if you have open systems disks installed, then z/VM can start to use those disks. And in fact, it doesn't require any account key data devices at all. So, this can help you reduce your total cost of ownership. And also, because the way we support this is through emulating SCSI disks as fixed block architecture devices, you get such capabilities as mini disk sharing, as well as mini disks themselves. So, you can partition a SCSI volume into multiple mini disks.

And we did this quite frankly because it was a way for us to leverage the existing support we had for FBA devices in Linux and in z/VM. So, for example, CMS and all of the other components of z/VM have no modifications in order to support SCSI disks using our FBA emulation. And while there is a little bit of native SCSI support code, there is not very much of it at all. So, it reduced our development expense quite significantly to take this approach.

Chart 10 describes the implementation. There's a new configuration file statement, the e-device statement, that defines an emulated device. Now the only kind of emulated device you can define at the moment is an emulated FBA device, but we anticipate that we may use this technology in future enhancements. So, you can define emulated devices, and you can delete them, of course, since they're totally made up – virtualized, if you will. And you can query them to see what devices are defined.

Currently, the largest device that we support, emulated FBA devices, is about 384 gigabytes. And one further restriction is that CP paging and spooling areas have to be within the first 64 gigabytes of a logical unit of a SCSI disk.

Now, there is some standalone support for SCSI disks in our standalone program loader, SALIPL. During CP IPL and shut down, we don't have the services of our full blown SCSI support, so we use standalone support. And also, when we're taking an abend dump with CP, we don't want to rely on the SCSI support that's in the nucleus. So, we have separate standalone support for writing a dump.

Chart 11 talks about another improvement we're introducing in z/VM 5.1, contiguous frame management improvement. The objective here is to reduce potential problems that we've seen in some of our stress environments in the laboratory here, associated with managing contiguous frames of storage. It's pretty easy for CP to allocate a single frame or even to allocate a lot of frames, as long as they don't have to be right next to each other. But when they have to be

contiguous, that presents something of a challenge to us. And sometimes, we can get into difficulties.

So, what we've done is improve some of the algorithms we use to allocate contiguous frames, and then we turned around and measured the stress workloads that showed some problems in the past to make sure that the benefits were delivered. So, we don't really think any customers have probably seen these problems. We're trying to prevent them from ever appearing in your environment.

Chart 12 details the various algorithmic changes that were made to support contiguous frame management better – such things as bringing in the rest of a block of pages, that is all of the pages except the one that was faulted on, above two gigabytes, rather than below two gigabytes – because of course, these contiguous frames are being allocated below the two gigabyte bar. And there are various other improvements that you can see here that, taken together, seem to have alleviated the problems that we were seeing in our test environments.

Chart 13 talks about our support for the PCIX cryptographic coprocessor. These coprocessor cards are available on the z890 and the z990. And the performance of these cards has been improved. And so, if you install these, you can support more clear key and secure key functions simultaneously. We support two kinds of access to these cards. One is through shared queues, and these are supported by Linux in its support of clear key encryption, for example, secure sockets, SSL.

We also support dedicated queues which can be used either for clear key or for secure key functions. And the secure key functions in particular are of interest to z/OS guests. So, we support both Linux and z/OS crypto functions. Essentially, what we have to do is recognize this new card type, because it is a new coprocessor, and reflect the card capabilities and various command responses.

Chart 14 talks about a function called HyperSwap. The objective here is to allow continuous operation of z/VM and its guests in a geographically dispersed parallel sysplex environment. That is a z/OS environment that has a multiple systems in geographically different locations which are usually separated by tens or hundreds of kilometers.

The approach we used here is to provide infrastructure in VM to dynamically swap devices. And what this does is take advantage of the fact that, in a geographically dispersed parallel sysplex, z/OS is establishing peer-to-peer remote copy relationships between local volumes and volumes at the remote site. And what this lets us do is; through the new HyperSwap command, it allows the Linux guest that's communicating with the GDPS infrastructure on z/OS to remove failing disks from the configuration and swap to the backup ones at the remote site.

Chart 15 shows this graphically. We have two sites in different locations, site one and site two. And in addition to having the ability for a site take over, for site two to take over if site one suffers some disaster, you see at the bottom we have this GDPS HyperSwap function that enables site one to start using the disks site two, should the disks at site one fail. And this is all managed automatically by the combination of z/OS, Linux and z/VM.

Chart 16 talks about a new feature called dynamic virtual machine timeout that lets a guest recover from a condition that it can't detect for itself. For example, a guest might be looping and may not include any logic to detect when it's looping. And so, it just sits there and loops forever. But this feature allows CP to detect when a guest has become non-responsive and allows it to specify what action is to occur in that eventuality.

So, the idea is kind of like the dead man switch concept, which of course is that you hold the switch closed – and actually, I think, it was related to engineers of steam engines. They would hold the switch closed and, if they fell dead for some reason, then they would let go of the switch,

the switch would open and the train would stop. So, the same idea is used here. A guest enables this timeout facility using a diagnose instruction and, at that time, defines the commands that are to be executed if the timeout actually occurs.

So, for example, the guest says, "Execute this log off command for me if a timeout occurs in 30 seconds." And then every 15 seconds say, the guests will reuse the diagnose to set it for another 30 seconds. As long as it keeps doing this, of course, the timeout never occurs and the log off command, or whatever it is, is not executed.

But if the timeout does occur, then the command is executed and the resources that the guest is using are free to be used by some other guests. So, chart 17 explains this is diagnose 288. It has three functions: initialize, change and cancel. And you specify an interval in seconds, and it must be at least 15 seconds. There's a SET TIMEBOMB command to let you turn off the effect of the timeout in case you're, for example, debugging a guest that's using this. You don't really want the time bomb blowing up in your face while you're trying to debug it. So, this allows the diagnose functions to act normally, but the timeout is actually suppressed. This state of the time bomb is reported by query set. And along the way, we made some enhancements to SET CONCEAL. So, it's now supported for a multi processor guest. It used to be a restriction that we've removed to facilitate this timeout.

Chart 18 talks about a new publication we're going to be providing called "Getting Started with Linux on zSeries." And the idea here is to make it easier to deploy Linux on zSeries by bringing together, actually, a lot of existing information, but putting it in one place so that it's easy to find and so that we can keep it up to date. So, this publication will be coming along with z/VM Version 5, Release 1 and expect it to be updated in subsequent releases and versions.

OK. That's it for the enhancements we made for Linux. But there are a lot of other things that we're planning to provide in Version 5.

Chart 19 lists some of the technology exploitation items – that is hardware exploitation items that we've included with Version 5, Release 1. The z990 and the z890 are introduced here, and there are some descriptions of their capabilities on charts 20, 21, 22 and 23. From z/VM's perspective, our support required us to test some capabilities that we introduced, actually in z/VM Version 4, Release 4. So, there's support for multiple logical channel subsystems. And the z990 and the z890 can both have multiple logical channel subsystems. We needed to test this new implementation of that function. Similarly, we have spanned internal and external channels and testing that support was also required.

Now, of course, we want to do this to make sure it's easy for customers to migrate to zSeries hardware, if you're not already on it, since Version 5 only supports zSeries processors.

Another thing related to the new hardware is on chart 25, support for larger single images. The objective here, of course, is to reduce the number of z/VM logical partitions you have to run. As our systems get larger and larger and can house more z/VM logical partitions, it becomes more complex to manage them. And when you divide resources among logical partitions, then there's less opportunity to share those resources most effectively. So, we wanted to address those problems.

And really, the work we had to do here was very straightforward from a technical point-of-view, but the heavy lifting was involved in evaluating the performance of these new larger configurations. We support up to 24 processors in a single image. And of course, this may not be appropriate for all kinds of workloads. So, we had to look at various workloads, and you can expect to see information about what we determined in the z/VM performance report, which will be available at the general availability of z/VM Version 5, Release 1.

In general, though, there are some guidelines relating to overhead that guests may experience on z/VM. If you have fewer guests and each guest is doing more work, then you'll probably see less overhead than if you have more lightweight guests.

One of the things that z/VM brings to the table, of course, is the ability to share resources, including sharing memory. And that's one of its wonderful capabilities, but you can go too far with this. So, just a warning here, if you overcommit storage to a great extent, you could cause additional overhead.

One of the new features on the z890 is an integrated communications controller. This is a new function of the OSA-Express adapter that is designed to reduce the total cost of ownership of zSeries, by not requiring you to have a 3174 or a 2074 associated with your z890.

So, there's a new CHIPID type OSC that we have to recognize and accept on the DEFINE CHIPID and DEFINE PATH commands, and report in response to a QUERY CHPID command. This adapter is also supported in z/VM 4.4. There's a PTF for the APAR listed on chart 26 at the bottom here that you can apply to your 4.4 system to get the support for an existing z/VM installation.

Chart 27 talks about the Enterprise Storage Server, the Shark model 750. This is a new low end Shark that reduces the cost of entry for adopting Shark technology. And essentially, we had to recognize the new device types, have device support, and enable the appropriate features. It supports all of the features of the more advanced models of the Shark, such as FlashCopy, and this comes in some smaller increments.

Chart 28 talks about some enhancements for capacity on demand. The zSeries processor lets you upgrade your processor capacity in a couple of different ways with on/off capacity on demand and capacity backup upgrade. And we support both of those with z/VM Version 5, Release 1.

And what we're trying to do is facilitate the automation of changes in processor capacity and processor capability. The approach we've taken is to recognize when these changes occur, report them through various mechanisms and to notify guests that are interested in these sorts of changes, as well.

So, we recognize when the processor capability changes – that is the processor speed may vary. And if it varies significantly, then this event is reported in the form of accounting – since you may want to charge a different price if your processor suddenly got faster. It produces a monitor record and displays a message to the operator. So, you could use an automation product to recognize that this event had occurred and take further actions if you wanted to. We also notify guests that are unable to receive these kinds of notifications.

If the processor configuration changes, if a new processor is added to the configuration or if one is removed from the configuration, then, again, we display an operator message and notify guests that are enabled to receive that information.

There's a new command QUERY CAPABILITY that reports the capability of the processor, and the capability is a number that's related to its performance. And we also changed the QUERY PROCESSORS command to tell you about processors that are now on standby mode – that have gone from reserve to standby mode, so you can determine which processors are available to be varied online.

Chart 30 talks about the 3592 model J70 tape controller, the model J1A tape drive. These are new technology, the controller and tape drive, that support twice as many FICON attachments as the 3590 model A60, up to eight SCON attachments and provide up to 1.5 times the throughput of the 3590 model A60.

z/VM supports these in all of the currently supported releases – 5.1, as well as 4.4., 4.3 and 3.1. And you can configure the device as either a 3490E or a 3590B. It's not a new device type from the viewpoint of the applications it supported. And there are PTFs required for CP and for DFSMS/VM shown at the bottom of chart 30.

Chart 31 lists a couple of enhancements to virtual networking and security, and we'll talk about each of these in turn. Our virtual switch has been very well received by customers. It provides a great way to provide network connectivity among guests and, also, to bridge that communication to a real network. And what we're doing with these enhancements is improving the availability characteristics of virtual switches, so that if there is a failure of a component, the recovery can be more timely and can be automatic rather than manual.

So, the way switches work, there's a switch controller – which is a virtual machine that sets up the environment that is needed for the real OSA device that's associated with the switch to operate. And then, it hands it over to VM. So, we allow multiple switch controllers so that there is no single point of failure. You can have a backup device for a switch – a backup OSA device with an alternate connection to the network in case there's a failure of the primary. We can now pre-condition those devices so that they're ready to jump into action immediately.

CP now periodically communicates with the switch controller to make sure that it's being responsive. And if it isn't, it automatically can switch in a backup controller. If there is a device failure, we can switch the queues that are being actively used to communicate with the device – to the backup device automatically so we don't have to stop and start the network connection.

And finally, we've limited the number of concurrent requests that we issued to the adapter so that we don't avoid it. When an adapter starts up, there's a lot of set up that has to be performed, and it's possible to overwhelm it with requests if they're sent too quickly.

So, there are some changes to the DEFINE/SET VSWITCH commands and corresponding changes to the QUERY of those entities.

Chart 33 shows a picture of the virtual switch. You can see we have seven VM guests labeled VM one, two, three, four, five, six, seven. And the virtual switch actually comprises a guest LAN and a real QDIO OSA device shown there in the middle of the chart, connected externally to an Ethernet LAN on which there are various other hosts. And you can see, there's a controller associated with that QDIO OSA device. And there's a dotted line between the controller and the OSA, because once it set it up, it then hands it off to CP and CP does the actual switching of traffic, either between the virtual machines or between the virtual machines in the network.

You also see a backup QDIO OSA device and an associated backup controller sitting in the environment here. And if the primary OSA device fails, then the backup OSA device is used automatically. And similarly if the primary controller fails, then the backup controller can take over the primary OSA device and provide whatever services it might need.

Chart 35 talks about some enhancements we've made to our authorization for virtual networking. Today, a guest LAN and a virtual switch have access control built into the control program itself. What we've done is optionally enabled an external security manager such as RACF for VM to control access to guest LANs, access to virtual switches and, also, to control the virtual LAN identifiers that a guest is authorized to use on one of these virtual networks.

So, we asked the ESM to authorize requests for access to guest LANs, virtual switches and to use VLAN identifiers. And only if the ESM doesn't support that function or if you don't have an ESM installed then is the CP access list used. So, there's the ESM interface that's activated during virtual network connection, which is either an explicit or an implicit COUPLE command to couple a virtual network interface card to a guest LAN or virtual switch. That is invoked to determine whether the guest is authorized to make the connection.

We've changed the way that revocation works. Today, when an authorization is revoked, the associated connection to the guest LAN or virtual switch is terminated immediately. That's no longer the case since, presumably, it was legitimate previously and you can take other action if you need to terminate the connection sooner than when user next requests access.

The CP access list associated with the guest LAN or virtual switch is still present. It can still be defined, but it may be ineffective if you have an ESM controlling access. So, we've added a new option to QUERY LAN and QUERY VSWITCH, the ACCESSLIST option, which only displays that list if you use that option.

And associated with this support are enhancements to RACF, the RACF Feature of z/VM, in z/VM 5.1 that exploit this new interface. So, you'll be able to take advantage of this feature right away.

Some real networking enhancements, as opposed to virtual networking enhancements, are listed on chart 37. The first of them is the introduction of support for IP Version 6. This is the first of what will probably be several steps to support IP v6 networks. And of course, we want to do this to provide address constraint relief. We're running out of 32-bit IP addresses. IP Version 6 provides 128-bit addresses. IP v6 also supports automatic configuration and introduces a variety of other improvements to the networking stack.

So, what we're supporting are separate Version 4 and Version 6 networks. And for Version 6 networks, we're enabling static routing, some basic applications and socket support.

Chart 39 talks about what exactly we've implemented. You can access IP Version 6 networks through OSA-Express adapters, through the QDIO adapters. And we provide static routing, as I mentioned. We also support router advertisements, so you can find a router on the network that

will be the master. And we support the trace route (TRACERTE), PING and IFCONFIG commands for Version 6. So, you can do some basic problem determination for connectivity and some basic configuration tasks.

IPv6 sockets are supported through the language environment and through open extensions callable (services) so you can write IPv6 sockets applications. But the Version 4 and Version 6 networks are treated independently, so there's no routing between the two types of networks. And in fact, there are separate HOME lists, separate filters, separate address translation tables, static routing tables and separate PORT lists for each type of network.

Chart 40 talks about a change we've made to simplify network configuration. We've added some more intelligence to our selection of default MTU sizes to try to reduce the probability that the MTU size will be specific incorrectly. So, you can explicitly say, "Give me the intelligent default," and z/VM TCP/IP will attempt to do that. So, you specify an MTU size as zero, and we select the default based on such things as the link type, the envelope size and the frame size.

And the MTU option has been added to the LINK statement for all interface types to make it easier to specify this, rather than having it buried in the gateway statement.

Chart 41 talks about a feature – it's not new in z/VM 5.1, but it's something that many customers don't know about, so I thought it was worth mentioning here. The TCP/IP server attempts to be autonomic in its control of other servers that are related to TCP/IP. So, it will restart a server when it detects that it's failed. But it's a little pathological about this, or it used to be in that it would keep restarting the server even if it was never going to succeed. Of course, it didn't know that. But if for example the server's a-disk was offline, then there would be no way it would be able to come out successfully. So, it would just fail repeatedly and we'd repeatedly restart it. So, now there's a MaxRestart statement that lets you specify the maximum number of times that a user will be restarted.

Chart 42 talks about a new threshold for SYN Flood. SYN Flood is a kind of denial of service attack that the z/VM TCP/IP detects and repels. Again, it was a little too aggressive – would sometimes report that an attack was occurring when it wasn't. So, we solved this problem, if you will, by letting you define when the threshold of such an attack should be deemed to be occurring. So, there's a PendingConnectionLimit statement that you can use to configure the maximum number of half-open connections – that is connections that have been offered but not consummated. And this is the signature of a SYN Flood, thousands of half-open connections until the stack that's being attacked finally fails.

So, when the limit is exceeded, whatever limit you set, we report that the attack is occurring and we select one of the half-open connections at random and drop it, so that we keep the number of half-open connections not above the limit that you specify.

Chart 43 lists a few other enhancements that we're making in z/VM Version 5, Release 1. First of all, you'll be able to install from DVD. This means there's one installation media, one disk that you can use not only to install z/VM, but also as an emergency recovery system. And while we introduce this primarily to support installation in a SCSI-only environment, in fact you can use this whether you're installing to SCSI or 3390 disks.

What we've done is create an IPLable starter system on the DVD so you use the load from server or CD-ROM function of the support element in the zSeries processor to IPL this DVD into the logical partition where z/VM will be installed. And we provide a support to load the installation files from this same DVD. So, you use the DVD drive and the host management console, or another FTP server that can be accessed by the zSeries support element.

Chart 45 gives a little more detail, this IPLable image is actually a self contained fully operationally z/VM system. It contains a CP nucleus, but in addition, it contains a RAMdisk.

We've borrowed some technology from Linux here. On that RAMdisk is a PARM disk, some of the system owned areas, the directory, warm start, checkpoint, paging and spooling, the operator's 191 disk and main's 190, 191, and 2CC disks. So, these are all of the things that you need to install to start the installation process.

So, you boot this from the support element using the load from CD ROM or server function, as I mentioned. And you can boot this anytime. So, if you accidentally do some damage to your system in such a way that it won't come up, you can bring up this system from the DVD and probably affect a repair.

Chart 46 talks about our virtual system management APIs. We've been rolling these out now for several releases. And these APIs provide support for system management functions. And we're adding some new APIs and providing some integration with our VM resource manager. So, we've added some volume management functions, image inquiry functions. We're supporting a new level of the remote procedure call interface, the RPC interface, and added support to manage the VM resource manager configuration file.

Chart 47 describes some enhancements to the Performance Toolkit. Performance Toolkit is the single performance management facility for z/VM now. As you'll see, this was a statement of direction that we made, and we're removing the Real Time Monitor and the VM Performance Reporting Facility and replacing them with a single comprehensive tool. Performance Toolkit is that tool.

We also, in addition to providing some new reporting functions, added some new reports and displays to exploit new performance measurements that are now being collected. In particular, we have new reports on SCSI performance and new reports for Linux guest. We've arranged to have Linux contribute performance information to the z/VM monitor data stream so that we can

actually get a picture of not only what VM thinks a guest is doing, but what the Linux guest itself believes is going on. So, you can drill down to the lower layers within Linux itself.

Chart 48 outlines some installation and service improvements. This is a continuing effort for us to simplify installation service and make it easier for you to stay current, to migrate to the latest release and version of z/VM. So, we've automated a few additional functions. And we've also taken the first step in providing a more stylized method for the migration from one release to another by separating IBM materials from customer materials. So, it's more clear cut where things go and who should be servicing them.

So, in particular, we've removed spool and page space from the system's residence volume since you may not wish replace your spool and page space when you install a new version of z/VM, but you do have to replace, of course, the system's residence. We've made it easier to display the RSU level and individual PTF levels and provided some better support for handling service memos. And we've also automated the local modification procedure. There's a new LOCALMOD command that you can use to facilitate handling local modifications you may have.

Chart 49 talks about some functions that we've removed. First is preferred guest support. Preferred guests were supported when z/VM ran in basic mode – that is when LPAR was not part of your configuration. But it turns out that most customers do run with logical partitioning. And in fact, the z890 and the z990 are LPAR only – that is you must run with LPAR on those systems.

And in an LPAR environment, z/VM doesn't support preferred guests. You can add a single V=R guest, but it doesn't benefit from the IO assist, the IO pass through capability that V=R was really designed to provide.

So, since there was less of a likelihood that preferred guest support would provide benefit and because most customers don't use preferred guest support and because we only support six

preferred guests – and that's probably not sufficient in a Linux environment on zSeries – we've decided to focus on improving performance for all guests by doing things for V=V, that is not preferred, or pageable guests.

We've also removed support for the ESA/390 Host architecture. What this means is that z/VM Version 5, runs on z/Architecture processors only. Now, you can still have an ESA/390 guest – CMS, for example, is through z/Architecture – but z/VM itself, the z/VM control program will always run in z/Architecture mode. And so, you must have z/Architecture processors.

A few more functions listed on the next chart, chart 50. SPTAPE command, as we promised, and a statement of direction has been superseded by SPXTAPE. That happened about 10 years ago. So, we finally removed SPTAPE from the product. ESA/370 Guest Architecture is not supported. It hasn't been supported for many years by our hardware, and certainly not supported by zSeries hardware. We do still support the 370 Accommodation features, so if you have 370 applications, they may very well run under z/VM with no problems.

As I mentioned, the RTM and PRF features have been replaced by the Performance Toolkit. We've also removed CMS support for Java and NetRexx. What we recommend you do is use Linux for zSeries if you want to run Java or NetRexx programs, and there's a URL you can refer to for more information about that.

Chart 51 lists some devices for which we've removed support – a number of older devices, disks, tapes and communication devices. You should really look in the general information manual for z/VM 5.1 when that's made available at general availability to see what the prices are and in fact support it and at what level.

Chart 52 lists some packaging changes. We no longer ship DFSMS/VM automatically. It's still available and it is a no charge feature, but you have to order it through the System Delivery Option.

Due to popular request, we've included the 3270 PC File Transfer product with the base z/VM. It's delivered as a sample, which is to say it's not supported, but it's part of the package because many customers have been looking for it. And it was no longer available.

We don't ship the Restricted source feature and the PL/X source. They're available for download if you want that material, but very few customers are interested in those materials.

We've removed the pre install of Tivoli storage manager. You can order it as a standalone product. It's not available through SDO. We really recommend, though, if you're interested in Tivoli Storage Manager that you look at running the server on Linux on zSeries, since that support is more current than the VM version.

We've also removed the National Language features for ISPF from the SDO. I don't know how many of you are actually interested in ISPF on VM. You can still get these features through the standalone ordering process if you are.

Chart 53 outlines some delivery changes. You can order VM and VM license products in the VM SDO over the Internet, using the ShopzSeries Web site. There's a portion of a screenshot shown here. And with ShopzSeries Release 8, which has been available since February of this year, licensed products for VM can be delivered via the Internet. VM itself still requires distribution media, but as I mentioned for Version 5, it will be a single DVD.

Chart 54 lists the statements of direction that we made with the announcement of z/VM 5.1 in April. We are currently in evaluation for Common Criteria certification at the EAL3+ level. Stay tuned as we go through the evaluation process, there will be more to say about that.

The second direction that we're taking is to support Layer 2 of the OSI reference model for OSA-Express, Virtual Switch, and Guest LANs. What this will do is enable network traffic to be routed using media access control addresses instead of IP addresses. Our OSA devices heretofore have used Layer 3, that is IP addressing, but will now use Layer 2 addressing. And this will enable a variety of applications on Linux such as IPX that were previously not supported. And I understand, this is also of interest to vendors of firewall products for Linux on zSeries. So, look for that to be delivered in the relatively near future.

z/VM version 5.1 is the last release where we will support some older terminal devices, the IBM 2741 and teletype devices or their equivalents. We don't think this is going to be a problem for anybody. If it is a problem for you, we'd certainly like to hear about it. And finally, we're attempting to withdraw the System Administration Facility and the Server Requestor Programming Interface.

Chart 55 talks about some of the areas in which we're focusing – or intend to be focusing in the future. First and foremost, we focus on virtualization for and synergy with Linux on zSeries. At the same time, we need to support the hardware so we focus on processor support, supporting new architecture, device support and networking, of course, especially virtually networking – because we think we have the best virtualization around, and we want to enhance our virtualization technology to make it as useful as we can.

And finally, as you may have noticed, we're continuing to spend on ease of use enhancements to make the system easier to configure, easier to manage and easier to install service and upgrade.

Susan Greenlee: Thank you, Romney. And to you, our audience, it should be evident that the evolution of IBM's premier and world class zSeries virtualization continues with the new enhancements we've made in z/VM Version 5, Release 1. And with the new engine base value unit pricing, you may see a price performance improvement in addition to a lower entry-point price.

That concludes our formal presentation for today. In just a moment, we'll begin our Q&A session. First, though, we'd like to ask you a few questions. Premiere announcer, can you tell our audience how to give us their feedback about today's session?

Operator: Thank you, Susan. To participate in our survey, please press the digit that corresponds to your choice, preceded by a star.

For our first question, please rate the overall effectiveness and value of today's call. If you found it very good, please press star, one. For good, press star, two. Press star, three, if you found it fair. Or star, four for poor.

Don't miss the LinuxWorld conference and expo August 2nd through 6th at the Moscone Center in San Francisco. IBM will once again be a platinum sponsor of the event and will have a large exhibit area, where you can see the IBM eServer family, including the new zSeries z890 server running a host of applications from IBM and other vendors on Linux. See Linux solutions in action that can help you run your on demand business and simplify your IT infrastructure.

As usual at LinuxWorld, IBM will host a customer day event, which will take place on Wednesday, August 4th. There will be informative presentations and discussions with IBM executives, business partners, developers and customers who are using Linux on IBM eServers.

Registration details for IBM customer day at LinuxWorld 2004, San Francisco, will be available in early July. Contact your IBM representative for details. For more information about the LinuxWorld conference and expo, visit the Web site at www.linuxworldexpo.com.

The second question in our survey is: where does your company stand with its implementation of Linux on zSeries? Press star, one if you have Linux and z/VM installed and an application in pilot or production. Press star, two if you have installed or are in the process of installing z/VM and Linux. Press star, three, if you have neither Linux nor z/VM installed.

Join us on June 29th at two p.m. eastern time for a teleconference on migrating workloads from Microsoft NT to Linux on zSeries. Migrating NT workloads to zSeries and Linux can help company's be flexible and adaptive to real time demand, simplify their environments and reduce total cost of ownership. You'll hear about how zSeries with Linux can support business applications in the areas of file and print serving, Web serving, database and e-mail functions. Call in to hear more about the advantages of migrating from NT to zSeries and Linux and to hear an example of a real life experience that highlights both the functional and financial advantages this migration offers you.

To register, please call 800-289-0583 or 719-386-0025 and specify confirmation code 455722. This information is also provided in today's presentation.

Our third question is: would you like to be contacted by IBM for answers to additional questions you have on today's topic? Please press star, one now if you would like to be contacted. Or press star, two if you do not want to be contacted.

Don't forget about SHARE. The next SHARE user driven training event and expo is August 15th through the 20th at the New York Hilton. SHARE provides an excellent technical program on a

broad range of IT topics, especially for zSeries and, in particular, for z/VM and Linux. For more information, please go to newyork.share.org.

And our final survey question today: which one of the following technical topics would you like to see presented as a one-hour teleconference next quarter? Press star, one for a technical topic of z/VM and Linux on zSeries installation. Press star, two for a technical topic about best security practices for z/VM and Linux on zSeries. Press star, three for a technical topic titled "Simplification of IT Infrastructures Using the New Communications Server for Linux on zSeries or SNA on Linux: Why Would I Do That?" Press star, four for a technical topic on Lotus Domino v6.5 for Linux on zSeries. Or press star, five for a technical topic on Samba.

And we'll pause for just a moment to compile your responses.

And that concludes our survey. Thank you, everyone, for your responses. At this point, we'd like to take your questions. To ask a question, please press star, one now on your telephone keypad. Again, that's star, one, if you do have a question today.

And while we give you a moment to enter the queue, I'd like to turn the conference back to Susan Greenlee.

Susan Greenlee: Thanks, Laura. Thank you very much for sticking with us and providing us with your feedback; we really appreciate it.

On our Q&A panel today, not only do we have Romney White, but Reed Mullen has also joined us from z/VM development in Endicott. Like Romney, Reed also speaks at user groups and at SHARE. Thanks for joining us today, Reed.

Reed Mullen: Thank you, Susan. Great to be here.

Susan Greenlee: Good. You know, I did a Google search on you Reed while we were sitting here listening to Romney. I see I can reach you at the first link that comes up in Google when I type in Reed Mullen – so if anyone has a question for you later.

Reed Mullen: OK. Great.

Susan Greenlee: Laura, do we have any questions?

Operator: Yes, we do. We'll go first to Rick Troth.

Rick Troth: Hey, Reed, Romney. Good afternoon.

Romney White: Hi, Rick.

Rick Troth: On chart 16, the virtual machine timeout, this is exciting new stuff. Let me ask if it's related to Linux's native watchdog timer support?

Romney White: Well, we certainly had watchdog in mind when we built this. And while we don't have any Linux exploitation that we talk about today, we certainly have some people in our Linux development lab who are very interested in exploiting this technology.

Rick Troth: I guess what I'm wondering is if we have to issue a CT command, or if the kernel performs these diagnoses on its own?

Romney White: Well, I think the idea would be to add support in the kernel to issue the diagnose upon request.

Rick Troth: OK. Looks like a neat idea. I hope we can make some good use of it. Thanks.

Romney White: Good.

Operator: We'll take our next question from Dennis Schaffer.

Dennis Schaffer: Hi, there. Can you describe some of the benefits of SCSI device ports as compared to the account key data device support?

Romney White: Well, let's see, first of all, SCSI devices can be considerably larger than account key data devices – at least today. And the 380 port gigabyte limit, I will tell you, is temporary. We intend to extend the size of lines that we support.

I think a second benefit – potential benefit of SCSI devices, depending on the infrastructure you have in place, the capacity for fabric is considerably better performance. As I say, this will depend on your configuration, but we've certainly seen some performance benefits from using SCSI.

And third, I think, the ability to share data, at least conceptually share data with open systems, the distributed systems is a larger significant benefit that you might be able to take advantage of.

Operator: Anything further, Dennis?

Dennis Schaffer: No. Thank you.

Operator: Thank you. We'll turn now to Richard Wiegand.

Richard Wiegand: Hello, this is Rich Wiegand. I was encouraged by all of the enhancements that you were putting in the VM for Linux. One of the things that goes along the lines of TCP/IP is secure socket layers. I was wondering if there's any future secure socket layers in VM support?

Romney White: Well, VM actually has support for SSL. In particular, we have an SSL server that has been in the product for several releases. It actually requires a Linux virtual machine in which to run. And what it does is automatically provide SSL encryption and authentication for servers that run on z/VM. So, the FTP server and the Telnet server, for example, can automatically be protected through SSL by implementing this server.

Richard Wiegand: Yes, and I guess I was aware of the (SuSI) SSL server. But what I was driving at is VM going to incorporate that in to their code in the future, rather than have an alternate operating system provide that functionality?

Romney White: Well, VM is all about running alternate operating systems like Linux and CMS and z/OS. So, it's not clear to me, although I'd be interested in your viewpoint on what some closer integration would provide.

Richard Wiegand: I guess it's just the mindset that there's an additional operating system that has to be purchased and supported in order to provide the functionality. That's where I was coming from.

Romney White: OK. I understand. Frankly, we'd like to provide you with a Linux, but we think that's better left to our distribution partners.

Richard Wiegand: So, the secure socket layers could only be that – that functionality could only be provided by a Linux partition? And it's not something that could be provided by VM?

Romney White: Well, it certainly could be provided by VM, but you know, we're talking about essentially migrating that function to CMS was what would be required. We certainly have no plans do that, but I understand your viewpoint and we can certainly take a look at it.

Richard Wiegand: OK. I was interested in what the future was for that – for that. And so, from what your statement is that VM will be looking for integrators to provide that functionality?

Romney White: Yes, I think that's fair.

Richard Wiegand: OK. Thank you

Operator: And I'd just like to make a quick reminder, if you do have a question today, please don't hesitate to press star, one now. And we'll take our next question from Adam Thornton.

Adam Thornton: Hi there. My question is the emulated FBA via SCSI support fully transparent to guests. That is, if I've got say some BSE guests with FBA disks, can I just put those disks on SCSI and have it work automatically?

Romney White: Yes, indeed. That was the whole idea, because we didn't want to have to make changes to CMS, for example, to enable it to use SCSI. So, it's fully transparent to guests.

Adam Thornton: OK. Great. Thank you.

Romney White: You're welcome.

Operator: We'll turn now to Paul Koniar.

Paul Koniar: Hello. Yes, I have a couple of questions. One of them is there any plans to integrate the performance toolkit functionality to more closely work with the IBM Tivoli suite? And the other question is are there any prop enhancement planned?

Romney White: Your first question about integrating performance toolkit with the Tivoli suite is a question for Tivoli, I think. We certainly provide the interfaces and the infrastructure that Tivoli would require to do such integration. I can understand why you would want to have that sort of function. I can carry the requirement forward to Tivoli, but it would be much more effective if you pass it through yourself, through your sales rep or through a user group share or our regional user group in – to let Tivoli know that this requirement isn't just something that we're making up here in VM land.

Paul Koniar: The idea being that kind of a, you know, common way to view all of the different platforms that were monitoring because performance toolkits are kind of a one-off right now.

Romney White: Sure. And of course, that is the Tivoli value proposition for their performance monitoring product. So, integrating VM with that makes a lot of sense to us. But you know, the only role we can play in that is supporting Tivoli's efforts. They have to initiate them.

Paul Koniar: OK.

Reed Mullen: Paul, I'll tell you – this is Reed. If I can just add to that. There is a Tivoli zSeries performance product that provides a portal for displaying a performance toolkit session. It doesn't provide the level of integration and commonality for all platforms that I'm sure you're looking for. But there is a minor piece of integration that's available. We do have their attention to some degree.

So, if you were successful in being able to get your request into other organizations through whatever channels you choose to pursue, it won't be unrecognized.

Paul Koniar: And if at the very least they're providing the single piece of glass...

Reed Mullen: Yes.

Paul Koniar: ... portal option to put up – put the window on the same screen so to speak...

Reed Mullen: Exactly.

Paul Koniar: ... big step in that direction.

Reed Mullen: Yes, and they do that with other products as well. So, the performance toolkit isn't unique in being – in having a portal.

Paul Koniar: Great.

Romney White: Right. That integration on the glass, as they say.

Paul Koniar: Yes.

Reed Mullen: Yes.

Romney White: As far as prop goes, I don't think we're looking particularly at making any enhancements, but we'd be interested to know what your requirements are.

Paul Koniar: Well, we're just getting in to the Linux arena now. And so, everything is new to us, including VM. And one of the things I'm – and I don't know what prop all gives us today, is how much capability do we have to monitor and automate the instances from prop – in terms of, you know, IPLing them – with your new diagnose tool, for instance, to be able to determine this one is dead, reboot it, re-IPL all of the instance.

Romney White: Right. Well I think there is a fair amount of capability already there. I'd encourage you to explore that. Perhaps commune with some of your fellow customers on the various Linux mailing lists. And if there is some need that you identify, some shortcoming, I encourage you to write to me or to Reed – since you can find him on Google – and let us know what your needs are.

Paul Koniar: Great. Thank you.

Susan Greenlee: Any other questions, Laura?

Operator: At this point, there are no further questions in the queue. Actually, I spoke too soon. We do have a question; we'll go to Dale Baker.

Dale Baker: Hi, this is Dale Baker with Abbott Laboratories. On foil number three, I want to make sure I understood that there has been a service discontinue now for z/VM Version 3, Release 1 on August of 2004 in about two months? Is that correct?

Romney White: No, that's not...

Reed Mullen: ...end of marketing. Yes, I'm sorry, Romney. That would be an end of marketing date. We have not yet announced an end of service date for Version 3.

Dale Baker: OK. I misread the red highlighting then. Thank you.

Reed Mullen: Yes.

Romney White: So, in other words, if you want to order it, order it before sometime in August – if you don't have Version 3 and you have a processor like a G3 9672 where you intend to run it.

Operator: And we'll turn back to Paul Koniar.

Paul Koniar: One last question – concerning the virtual switch support, some of the testing we've been doing has come up with some oddities in the way information is returned through Linux query commands. We're running – we're trying to run like three virtual LANs. And I'm wondering if some of this virtual switch support will help clear up some of those issues, including the fact that, you know, you have a standard – when you're using virtual LANs, you have your standard connection like your ETH0. And then, you set up like a VLAN definition.

And when you go to query to find out where the traffic is coming out, it doesn't show up under the VLAN one like you expect; it shows up under ETH0, which has caused some issues with some of the vendor products that we're working with.

Romney White: I guess if you think you have an issue with VM support, then calling the IBM support center would be the best way to resolve that. The virtual switch is certainly easier to deal with, I think. And it also offers a better performance for Linux guests than using a guest LAN, because you can eliminate the need for a router.

But I'm not aware specifically of any problems with Linux, but that doesn't mean there aren't any. If you think it's a VM thing, by all means call our support folks.

Paul Koniar: Great. Thanks.

Operator: And just a final reminder, if you do have a question today, please signal us by pressing star, one.

Susan, it appears there are no further questions today. So, I'll go ahead and turn the call back over to you.

Susan Greenlee: OK. Great. Thank you, Laura.

I guess we should close our session now. If you have a question and you weren't able to ask it, or if you think of one at a later time, you can send a note to Romney. His e-mail address is on the front of the presentation. Or you can send me an e-mail, it's SGreenle – sgreenle – @us.ibm.com. Or you know how to get to Reed Mullen.

And let me just point you to the last page of the presentation. It supplies some pertinent information about our call. The first link there is to the replay. If you want to point some of your peers to the MP3 that will be available for download, it will be at the same sight where you got the presentation today. And then, if you're looking for documentation on z/VM 5.1, you can check out that second link there. And also, you may find the software group matrix link for the middleware products that run on Linux on zSeries. That can be very helpful.

On behalf of Romney and Reed and myself, thank you for attending our session. And I look forward to joining you again next quarter. Thanks.

Operator: With that, we'll conclude today's program. Thank you, everyone, for your participation.

END