April, 2008

Determining which Solutions are the Best Fit for Linux on System z Workloads

Bill Reeder, IBM

Table of Contents

The purpose of this paper is to discuss the current state of the art as to hosting applications on the Linux® on IBM System z™ platform, and to help in the decision process of what to host where. In order to help gain clarification in helping customers decide what could or should be examined for migration to Linux on System z, applications/functions are divided into three broad categories, and the descriptions of the applications attempt to qualify the migration rational deeper. The three categories are:

- ***Best Fit*** *solutions for Linux on System z.*
- ***Good Fit*** *solutions for Linux that can be deployed on System z, IBM System x™, IBM System i™ or IBM System p™ platforms. This section will elaborate on some of the decision criteria involved in platform selection.*
- ***Challenging Applications*** *are solutions that* ***generally*** *best fit on architectures other than the System z platform.*

Best Fit Applications

In general these are applications that leverage the System z platform classic strengths as well as being the source of data and application information. The classic strengths of the System z platform that are of most importance for best fit applications are high availability, high I/O bandwidth capabilities, the flexibility to run disparate workloads concurrently, and excellent disaster recovery characteristics. Applications that are used to communicate directly with legacy mainframe applications are able to leverage architectural advantages of the System z platform. As a result we often find that the combination of these applications on a single platform far outperforms the predicted sizing models of the individual application components. In the newest System z platform the IBM System z10™ Enterprise Class (z10 EC™) with new high clock speeds in excess of 4 GHz, and the newer 10 Gb/s OSA cards and 6 GB/s HiperSockets™ speeds, the number of options that put applications into this category have increased. Examples of these are MQSeries®, DB2 Connect™ and IBI Web Focus from the list below.

- *WebSphere® MQSeries; MQSeries is almost always deployed on distributed platforms, yet in most situations MQ is being used in conjunction with legacy mainframe applications. Moving MQ queues from the distributed environments closer to source data and leveraging HiperSockets, gives the net effect of simplifying the network, improving processing throughput, and the overall management of the data flow from end-to-end, reducing the number of places one needs to go in troubleshooting a problem. Additional savings and performance benefits can be realized if the application being communicated to with MQ is also located on the same platform, such as WebSphere, and Apache applications.*

- *DB2 Connect; just as MQSeries, has the single purpose of connecting applications to data. In this case the data is stored in DB2® databases running under z/OS® or z/VSE™. Moving DB2 Connect from the distributed environments to Linux on System z, brings the connectors closer to the source data, and could easily experience performance improvements by leveraging HiperSockets. This has the net effect of simplifying the network and improving processing throughput, it also improves the overall management of the data flow from end-to-end, thus reducing the number of places one needs to go in order to troubleshoot difficulties. Additional savings and performance benefits can be realized if the application utilizing DB2 Connect is also located on the same platform, such as WebSphere or Apache applications.*
- *CICS® Transaction Gateway; Transaction gateway is generally used to Web enable traditional CICS applications. Like DB2 Connect and MQ above a great deal of savings and performance benefits can be realized if the application being communicated to with Transaction Gateway is also located on the same platform, such as WebSphere, and Apache applications. Other benefits are leveraging existing business logic while adding Web browser front ends along with adding required end-user functionality.*
- *IMS™ Connect for Java® can connect to IMS Connect running under z/OS to access existing IMS applications. Like DB2 Connect, Transaction Gateway and MQ above a great deal of savings and performance benefits can be realized if the application being communicated to with IMS Connect is also located on the same platform, such as WebSphere, and Apache applications. Other benefits are leveraging existing business logic while adding Web browser front ends along with adding required end-user functionality.*
- *Oracle Database is a powerful candidate to look at re-hosting. The current GA releases are 9i, and 10g both r1 and r2. Oracle 9i on System z release is a 31-bit version meaning those databases should be kept to less than 300 GB and the System Global Area (SGA) size less than 1.5 GB. Realistically this fits many applications that are developed by customers, as well as ISV applications. Oracle supports most all of their applications in a split tier environment with the database on the System z server and the application server hosted elsewhere. The 10g version is 64-bit only, and the SGA and database size limits do not apply with regard to memory. Oracle database on System z also support Oracle Clusters (Oracle RAC). Again the advantages to Oracle RAC on System z Linux is the low latency within any one platform, and with the newer 10 Gb/S combined with HiperSockets puts the scalability and performance of an Oracle RAC cluster customer in a competitive position over other most other architectures.*
- *SAP has been available on the System z platform with the Database DB2 on z/OS and the application servers available on Linux. This solution has a distinct advantage of keeping the people costs for systems administration under control and also limiting the environmental impacts such as excess use of floor space, power, and power for cooling. SAP solutions using a System z framework of z/OS and Linux on System z have been very successfully implemented in SMB, Government, and very large accounts. Often this solution is also accommodated by a subset of the servers also running on IBM Power Architecture® to create an optimal and well balanced high performance and cost effective model.*

- *WebSphere and Java applications development platform; While the developer generally develops on their own platform, the testing and modifications are generally done in a server environment. There are several different ways that customers can and do leverage the Linux environment on System z to decrease the time to market with applications. Developers can be given multiple virtual servers to perform interactive testing while troubleshooting or enhancing the application. Using the z/VM® mini-disk images with several Linux guests mounting to the same image, the customer can guarantee that the developers are all using the same versions of the Kernel, and keep problems from creeping into the development process. One other big advantage being seen by some customers is centralization of the development servers in a global development process. Simply put, the virtual machines are available around the clock. Similarly customers whose bulk CPU intensive work is at a fixed time of the day can easily utilize the non busy CPU times of the day for development work.*
- *Rapid Application deployment platform is accomplished by rapidly deploying and purposing Linux Virtual machines. This can satisfy a business requirement much more quickly than the normal procurement process. This capability allows customers to deploy servers in minutes that can either be individually customized or can be clones of existing servers. Along the same line of thought, should a new application be discarded, for some business reason, then the virtual machine can simply be discarded, as well. The resources are merely returned to the general pool for re-use as needed. In the world of fixed machines the business has much more difficulty finding a use for the unused test prototype machines. With the higher clock speed and the improved decimal floating point in hardware capabilities, of the new z10 EC processor, applications which after the development phase might have been moved to distributed architectures can take advantage of improved local machine performance again co-hosting the DB2 database that would support the WebSphere application. Customers can eliminate the DMZ layer between the application and the database, the customer should also leverage HiperSockets, then the transaction throughput capabilities of this kind of workload on System z is an excellent price and performance solution.*
- *Domino™ solutions have become an incredible candidate for migrations onto System z and Linux. An example of consolidation that is presently being done is within IBM's own internal IT. Some of the easiest candidates to move are Domino applications. The reason behind this has a lot to do with the design point of Domino as an application server. Domino allows for a reasonable seamless migration of applications between platform architectures. This design point of Domino is such that movement of both Domino email and applications between architectures is done very seamlessly, making Domino an excellent environment for allowing customers a high degree of flexibility in where to deploy applications.*
- *Network Infrastructure, FTP, NFS, DNS etc.., are well served on the Linux on System z server. Generally the workloads represented by these services are minimal, yet they are critical to the business. The main advantage hosting these services from Linux on System z are the availability of the hardware, disaster recovery capabilities, and the ability to host these small but critical workloads. Since a great deal of network traffic is generally generated between the legacy mainframe data and the FTP or NFS servers, if the FTP, or NFS data server is hosted on the same machine as the data source (Mainframe) and HiperSockets are used, then not only is the network traffic greatly reduced, but the batch processing window for that data can also be reduced.*

- *Applications which require easy and rapid failover are excellent candidates for deployment on Linux on System z. This functionality can be delivered using a SYSPLEX or Geographically Dispersed Parallel Sysplex™ (GDPS®). This is a very new and important feature that takes advantage of the SYSPLEX being managed in a z/OS environment and using storage mirroring of XRC or PPRC. Linux guests can leverage communications to create an asynchronous failover guest in a geographically dispersed customer environment. The methods for using GDPS involve the failover structures being under control of z/OS, and are in essence an asynchronous communication and synchronization of the disk system supporting the applications. Should a failover occur the Linux guest would fail-over along with the rest of the GDPS system.*
- *Communications Server fits extremely well and helps to solve a problem with customers that are still using 3745 and 3746 front end processors. As of May 2004, we are able to support TN3270, hosting as well as LU 0, 1, 2 and 6.2 SNA communications. The current release will allow the customer to install the base Communications Server code and gain valuable experience with the product. In 2005 much of the NCP, SNI and partner communications code that helps customers retire their old front end processors became available. One of the more important TCO advantages to moving this functionality onto Communications Server, is that the leased line costs for SNA traffic are up to 7 times more expensive than IP traffic. It is also important to realize that IBM support for 3745 and 3746 front end processors ends in 2010.*
- *Communication Controller for Linux (CCL), allows customers to migrate workload from Front End Processors (FEPs) such as the 3745/3746 and 2216. Specifically the customer can generate the Network Control Program (NCP) and run the NCP workload directly on the System z processor. SNI communications is also accommodated by CCL. At this point depending on what FEP functions the customer has they can generally reduce the number for FEPs using CCL as a consolidation vehicle. In some cases where all of the supported functions such as SNI are supported, the FEP can be removed entirely. CCL only runs on Linux on System z.*
- *LDAP security services fit very well running under Linux on System z. There are both open source OPEN LDAP products as well as commercial products like Tivoli® Directory Server, Tivoli Directory Integrator and Tivoli Access Manager (TAM). Using the System z architecture customers can create a robust 24x7 LDAP infrastructure to keep applications authenticating to LDAP working and available. Security can also be easily tied from a User ID and password standpoint to security services running under z/OS, z/VSE and z/VM.*
- *IBI WebFOCUS is a Meta data warehouse type of solution and is a great example of leveraging a multi-layer application on platform as a consolidation model that has the ability to satisfy customer demands with the System z platform. In customer benchmarks with Web Focus and the database running on UNIX® servers compared to Web Focus running on Linux on System z communicating over HiperSockets to the database running DB2 on z/OS, performance based on throughput for like priced serving was found to often be more than 5 times the performance. This was not a result that would have been anticipated using typical piece parts sizing models.*

- *Communigate Pro is a Voice over IP (VoIP) solution, VoIP represents a new area for the System z platform. Benchmarks were run in 2007 that showed that the System z platform was able to host VoIP services for 25 Million users, over 5 Million connections simultaneously. Previous benchmarks on distributed hardware solutions were not able to scale beyond 1 Million users. The principle reason for the large scaling of VoIP solutions hosted on the System z platform is that there is generally a great deal of communications between the various servers in a VoIP solution, once again the HiperSockets communications has provided an environment, that is not achievable in other technologies.*

Good Fit Applications

Applications and functions in this category are those that leverage the strengths of the System z platform, as well as strengths and business requirements of the customer. Applications in this category, while they can fit well running on Linux on System z, will also fit and run very well on other architectures such as System x and System p, using both AIX® and Linux. There are many variables that customers need to evaluate in order to determine which platform or group of platforms is the best fit. There is no ONE answer and in many cases a Hybrid Server environment may be the best fit. The platform of choice decision process needs to look at items such as:

- *Total Cost of Ownership (TCO). It is very important that the TCO model consider ALL the variables, not just hardware and software. There needs to be consideration for operations cost, environmental factors, and for many businesses the cost of interrupted operations for whatever reason. There are other factors that weigh into the total cost equation as well such as training, and charge back models. One element that is frequently missed in this arena is the inclusion of* ***ALL*** *of the systems needed for the complete solution, including backup, recovery, education, and test systems. All of these servers need to be included to get the true picture of what the enterprise will experience in making platform architectural choices. It is important to weigh the Total Cost and not just hardware and software costs when performing a TCO analysis.*
- *Experience and staffing considerations, can be an important part of the TCO decision. Consider a customer migrating from SUN making a decision to move to Linux or AIX. From a systems administration standpoint, the amount of training and change is about the same. If moving to Linux running under z/VM is in the mix, then there also needs to be some consideration for those skills as well. Frequently those skills exist locally, and experience has shown that the disciplines in place in the mainframe environment can be a determining factor.*
- *Politics within the customer site can never be underestimated. Most customers have structural pillars such as Windows®, UNIX and Mainframe. Each of those pillars is often competing within the overall organization for the same workload, and are each more than capable of making a very compelling case. As much as possible I try to advocate "what is best for the customer" from a cost, performance, and availability perspective. In an ideal situation this should eliminate any bias toward an individual architecture.*

- *Looking again at proximity to the core data of a distributed application should also be part of the evaluation. Frequently the data feeding the distributed application originates on the legacy mainframe. If one examines the overall dataflow, and the complications, then ask the question "is that still necessary, or can there be savings in batch processing, management costs, disaster recovery, that can be realized by co-hosting a distributed application on Linux on System z, within the same architecture of z/OS or z/VSE where the data that is running the business originates".*
- *Applications that have High I/O are generally good to examine running on Linux on System z since the processing of I/O workloads is offloaded from the main processor by the service assist processor. This is often argued over but is not to be overlooked. Keep in mind also that unused cycles such as those from test and development machines can be fully available to the production system.*
- *Hybrid applications environments will leverage the availability and performance requirements that best suit the business situation. It is important to examine the end-to-end transaction, and not just the pieces and parts. While there are many examples one is provided here. Imagine that you have deployed your database using z/OS DB2 connected to CICS in a Parallel Sysplex® cluster for the purpose of achieving high availability and substantial ability to scale. The customer then begins a rapid application development cycle to get a WebSphere deployment model moving. The development process is enhanced by the availability of a large server pool to the developers. The actual deployment model then for this is being done in clustered System p servers. The DMZ security and flat Web serving (HTML only) front end is being handled by some xSeries® Blade servers, that also serve the purpose of being able to handle the outgoing Fax requirements for this application. All of these environments must be considered holistically to understand the end-to-end application.*
  - *DB2 Universal Database™ (DB2 UDB) is just now being released as the 64-bit version running on Linux on System z. At present most of the UDB that is running is done supporting WebSphere applications. It is expected that growth here will continue as more ISVs are beginning to support UDB running under Linux on System z as a supported platform. If quality of service requires functionality under z/OS, such as DB2 data sharing and Parallel Sysplex exploitation, then consider keeping or moving the database to DB2 on z/OS.*
  - *Informix, (IDS) has been 64-bit capable for quite some time and is an excellent place to move customers that have been running Informix on other platforms. Customers considering running Informix on Linux are generally trying to leverage the uptime and disaster recovery capabilities of the mainframe environment.*
  - *Apache Web serving has some other variables that need to be looked at. First and most frequently is what product is used in the current Web serving environment. If it is Apache, migrations are very straight forward, and a normal architectural assessment and business cases apply. Migrations of classic PHP, My SQL Apache Web hosting environments is very straight forward, and about as close to drag and drop as migration gets. The second condition is if the current server is Microsoft® IIS server, then there are a few more things that need to be examined. If the code is running Active Server Pages (ASP) then that code needs to be dealt with. There are several applications that can help customers convert their Microsoft applications developed with Visual Studio to run purely under Java. The underlying application programmer interfaces (APIs) from the Microsoft .NET framework have been written by the Open Source community into a framework known as the MONO project. Taking the underlying*

APIs from a Visual Studio developed application and recompiling the code with software from companies such as Mainsoft allows those applications to be compiled into a J2EE environment, allowing them to be then easily hosted in a J2EE framework such as WebSphere, and many architectures

- *SAMBA can also be a very good fit for running on the System z platform. SAMBA services are fairly easy to implement and we have customers with more than 2 TB of data being served. SAMBA can authenticate to Microsoft's Active Directory (AD), or LDAP, or to the Linux password file. Additionally SAMBA can be used as just a simple access method for people like developers to move files into a server directory like in APACHE. The point being that some customers will say they don't want to run SAMBA because they have "made a Microsoft decision" and that Microsoft structure will be the source for file and print serving. SAMBA can cooperate in that structure, and if for political/business reasons they don't want to use it for file serving, they can, and really should still use it as an easy file access method where appropriate.*

Challenging Applications

Applications that may be challenging are those that though they may run on Linux on System z present issues that have either a business or technical reasons for examining the workload further and seriously considering another architectural solution.

- *ISV and IBM applications that have not yet ported their application to run on Linux on System z. If the application is not available on the System z platform, and the ISV has not indicated any interest in porting to the System z platform, then it is only in those rarest cases that we can meet the client's need by pursuing an application port. I strongly recommend working with the ISV and making the request for the future, but at the same time pick another project to migrate an application to the System z platform.*
- *Applications that by design run at **VERY** High sustained utilization which I will define here as >98%. These applications are of a type like geological mapping, animation rendering and the like. They are designed to run at a sustained >95% often at 100%. I would strongly recommend putting these onto distributed systems like Blade centers with Intel® or JS20 blades and optimize your investment and throughput.*
- *Stand-Alone single applications as the only Linux on System z applications, with no plans for further utilization. If a customer is trying to just move a single application that may occupy a few servers, and really has no further plans, then Linux on System z may not be the best fit. One main reason is that the customer may be buying more capacity and infrastructure than is needed. Simply put the more that is running on Linux on System z the better the value proposition as many of the infrastructure costs and operations costs are better amortized over the larger scope of projects.*
- *Applications that are too internally sensitive to try and move. The organizational structure of the company and IT organization as well as the commitment by many people to the existing way of hosting the application can be too much to overcome. Often in attempting to migrate these kinds of applications a continual stream of roadblocks will appear from those resisting the change.*

Finding Supported Applications

The URL below lists the current known supported applications that are supported and run under Linux on System z. There may be ISVs not on this list that may be in a beta release or are trying to build their business case. Without a doubt it is the customer telling the ISV that they intend to run their application on Linux on System z that will offer the biggest chance to make that happen. Communicating back to us and the ISV team also helps a great deal. I can be reached at breeder@us.ibm.com and will forward requests to the appropriate team member, or contact the ISV directly.
http://www-1.**ibm.com**/servers/eserver/zseries/os/linux/apps/all.html

This URL is a link to the IBM software running and supported to run under Linux on System z.
http://www-1.**ibm.com**/servers/eserver/zseries/os/linux/software.html

Conclusion

Linux on System z is an attractive solution for many business problems. The decision points that would help a customer to decide which application is the best fit have been a work in progress. The availability of applications along with some great product enhancement both in hardware and code of the System z machines has made Linux on System z an ever improving choice. This paper has been designed as an initial guide and is in no way the definitive answer. More importantly the material is constantly evolving, and really represents an overview to this point. Please engage the IBM team or the IBM business partner team to help validate you proposed solution and work with you to find the best fit.

IBM Corporation
New Orchard Rd.
Armonk, NY 10504
Printed in the United States of America, 04/2008

ZSW03037-USEN-00