Best Bank In Town
Since 1872

New Workloads for System Z

March 15, 2010

Outline

BB&T Best Bank In Town Since 1872

1. BB&T
2. Strengths of the System Z Platform
3. Why consider zLinux?
    1. Server Consolidation isn't new – why on System Z?
    2. Cost model
    3. Environmentals
    4. Software licenses
    5. Disaster Recovery
4. Workloads
    1. Apache, WAS, DB2, Oracle
    2. Characteristics
5. Staff considerations
6. Weaknesses of the System Z Platform (vs. P-Series and x86)
7. Where do you go from here?

1. BB&T

BB&T Best Bank In Town Since 1872

Building on a tradition of excellence in community banking that stretches back to 1872, BB&T continues to offer clients a complete range of financial services including banking, lending, insurance, trust, and wealth management solutions.

- $165.8 Billion in assets
- 10th Largest Financial Holding Company in the United States
- Over 1,800 branches in 12 states (plus Washington, D.C.)
- A little over 32,000 employees
- Strong market share in our footprint (#1 in West Virginia, #2 in North Carolina, #3 in Virginia and South Carolina, #4 in Alabama and Kentucky, #5 in Georgia and Florida, #6 in Tennessee and Maryland, and #7 in Washington, D.C.
- BB&T Insurance Services - #7 Retail Insurance Broker, #1 in Carolinas and Virginia, 8th largest brokerage worldwide
- Scott & Stringfellow Retail Brokerage Services
- BB&T Investment Services – over $4.6 Billion invested

2. Strengths of the System Z Platform

BB&T Best Bank In Town Since 1872

2. Hardware Overview

BB&T Best Bank In Town Since 1872

- The whole frame is called a CEC ("Central Electronic Complex")

- Processors and memory are added in "books"

- Books are hot pluggable and hot swappable. Hot swapping assumes you have enough spare processors and memory on other books to run the workload.

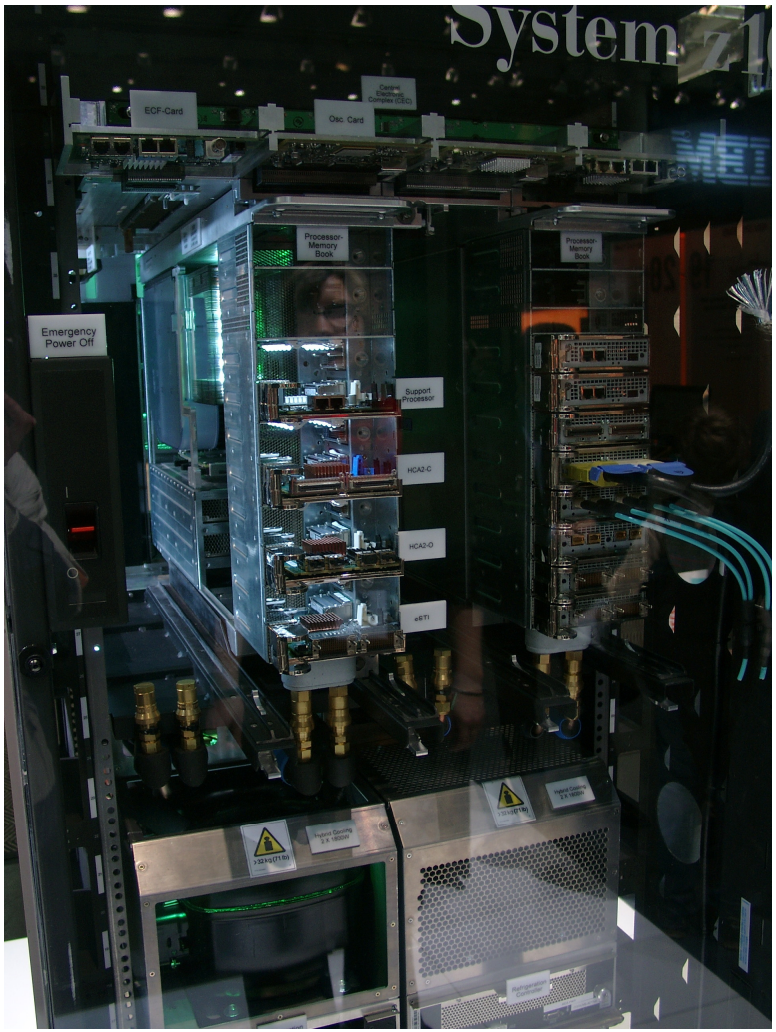

2.1 Books, Processors and Memory

BB&T Best Bank In Town Since 1872

Model	# cores	Available Processor Cores	System Assist Processors	Spare Processor Cores	Max Memory
2097-E12	17	12	3	2	352 gig
2097-E26	34	26	6	2	752 gig
2097-E40	51	40	9	2	1136 gig
2097-E56	68	56	10	2	1520 gig
2097-E64	77	64	11	2	1520 gig

2.2 I/O Capabilities

BB&T Best Bank In Town Since 1872

- One, two or three I/O cages
- Each cage has 28 I/O card slots
- Max of 3 cages X 28 slots = 84 I/O cards
- 84 cards x 4 ports = 336 8 Gb/second ports (maximum for FICON)
- Hot pluggable, hot swappable

- Lots of card choices

2.3 What does all that mean?

BB&T Best Bank In Town Since 1872

- Better at doing I/O than any other computer made. Ever. Not kidding.
- Extremely good at context switching
- Not so good at analytical compute intensive applications – put them on a RISC platform (P-Series).
- The most reliable platform for running mission critical applications (zOS).
- zVM is not too far behind – very stable, but not as robust as zOS.

3. Why consider zLinux?

BB&T Best Bank In Town Since 1872

- Server consolidation is not new. VM is old, VMWare has been around a while and is very prevalent, Power Virtualization looking more and more like a mainframe every day
- Other virtualization tools and platforms are gaining momentum (Xen)

- zLinux makes sense where:
  1. You already have a significant investment in System Z, particularly if you have spare engines in your mainframe that are not “turned on”
  2. Already have experience in the organization dealing with virtualization
  3. Have existing deployments of middleware that is fully supported and ports well to zLinux (WebSphere, Oracle, DB2, MQ)
  4. z9’s are ok, z10’s are MUCH better. Processor speed is a huge advantage with the z10. z9s will limit your choices.
  5. Are already spending significant $ on new server hardware every year

3.2 Cost Model

BB&T Best Bank In Town Since 1872

- The model on the next page is the basis for BB&T's cost model.
- The prices are either list price, or wild guesses, and all have 0% discount applied.
- Assumes 5 year (60 month) depreciation
- Assumes annual maintenance is 10% of the purchase cost
- SAN cost is $1.25 per gigabyte per month – sample cost only
- Assumes 1 FTE for every 5 IFLs to maintain zVM and zLinux environment. This is very, very conservative – likely nowhere close to that, but it's a starting point.
- Assumes 16 GB of memory per IFL. Assume you can over subscribe.
- Assumes a ratio of 15 servers per IFL for production and 30:1 for test, and assumes test does not need DR

3.2 Cost Model

BB&T Best Bank In Town Since 1872

zLinux cost model

	List	Discount %	BB&T Price	Depreciation Months	Monthly Depreciation	Maintenance - List	Maintenance Discount %	Maintenance - Discounted	Maintenance - Monthly	Total Monthly - Production	Total Monthly - Test
Hardware											
IFL	75000	0%	75000	60	1250	7500	0%	7500	625	1875	1875
DR IFL	75000	0%	75000	60	1250	7500	0%	7500	625	1875	
Memory - 16 Gig	36000	0%	36000	60	600	720	0%	720	60	660	660
DR Memory - 16GB	36000	0%	36000	60	600	720	0%	720	60	660	
Software											
zVM	25000	0%	25000	60	417	5630	0%	5630	469	886	886
zVM RACF	5000	0%	5000	60	83	750	0%	750	63	146	146
Other Linux tools	10000	0%	10000	60	167	750	0%	750	63	229	229
zLinux (RHEL)	20000	0%	20000	60	333	500	0%	500	42	375	375
DirMaint	1000	0%	1000	60	17	100	0%	100	8	25	25
Monitoring tool	20000	0%	20000	60	333	2000	0%	2000	167	500	500
Perf Toolkit	3500	0%	3500	60	58	350	0%	350	29	88	88
WebSphere	20000	0%	20000	60	333	2000	0%	2000	167	500	500
DB2	40000	0%	40000	60	667	4000	0%	4000	333	1000	1000
Staff											
Staff (0.2 FTE/IFL)	100000	80%	20000	12	1667		0%	0	0	1667	1667
										10485	7950

	$/GB/month	# GB per zLinux - prod	# GB per zLinux - test		
Disk space - 40GB per instance	$1.25	40	40	50	50
# zLinux/IFL - Prod	15				
# zLinux/IFL - Test	30				
Monthly Chargeback - Prod	749				
Monthly Chargeback - Test	315				

3.3 Environmentals

BB&T Best Bank In Town Since 1872

- Much more difficult to quantify
- Not included in my cost model

- Power, cooling, floor space in the datacenter
- Network switch ports
- Cables (power, network, SAN connnections, etc.)
- Racks

- All of these things are expensive, and will lower the TCO of zLinux

3.4 Software Licenses

BB&T Best Bank In Town Since 1872

- Example - WebSphere:
- 3 physical IFLs
- 3 copies of WebSphere
- On a z10, 120 PVUs per IFL, 3 x 120 = 360 PVUs
- 360 PVUs x $150 = **$54,000**

- On physical hardware:
- 15:1 consolidation ratio of servers to IFLs
- 45 processors, assume 100 PVUs each
- 45 * 100 * $150 = $675,000
- It gets worse – assume at least a dual core processor on each physical server on distributed hardware:
- 90 processors * 100 * $150 = **$1,350,000**

- And that's just for WebSphere!!!!

3.6 Disaster Recovery

BB&T Best Bank In Town Since 1872

- On any virtualized platform, DR is much easier than with dedicated hard drives.
- No internal hard drive – all data on the SAN (or your mainframe disk)
- Entire server is restored, along with it's configuration files

3.6 Disaster Recovery

BB&T Best Bank In Town Since 1872

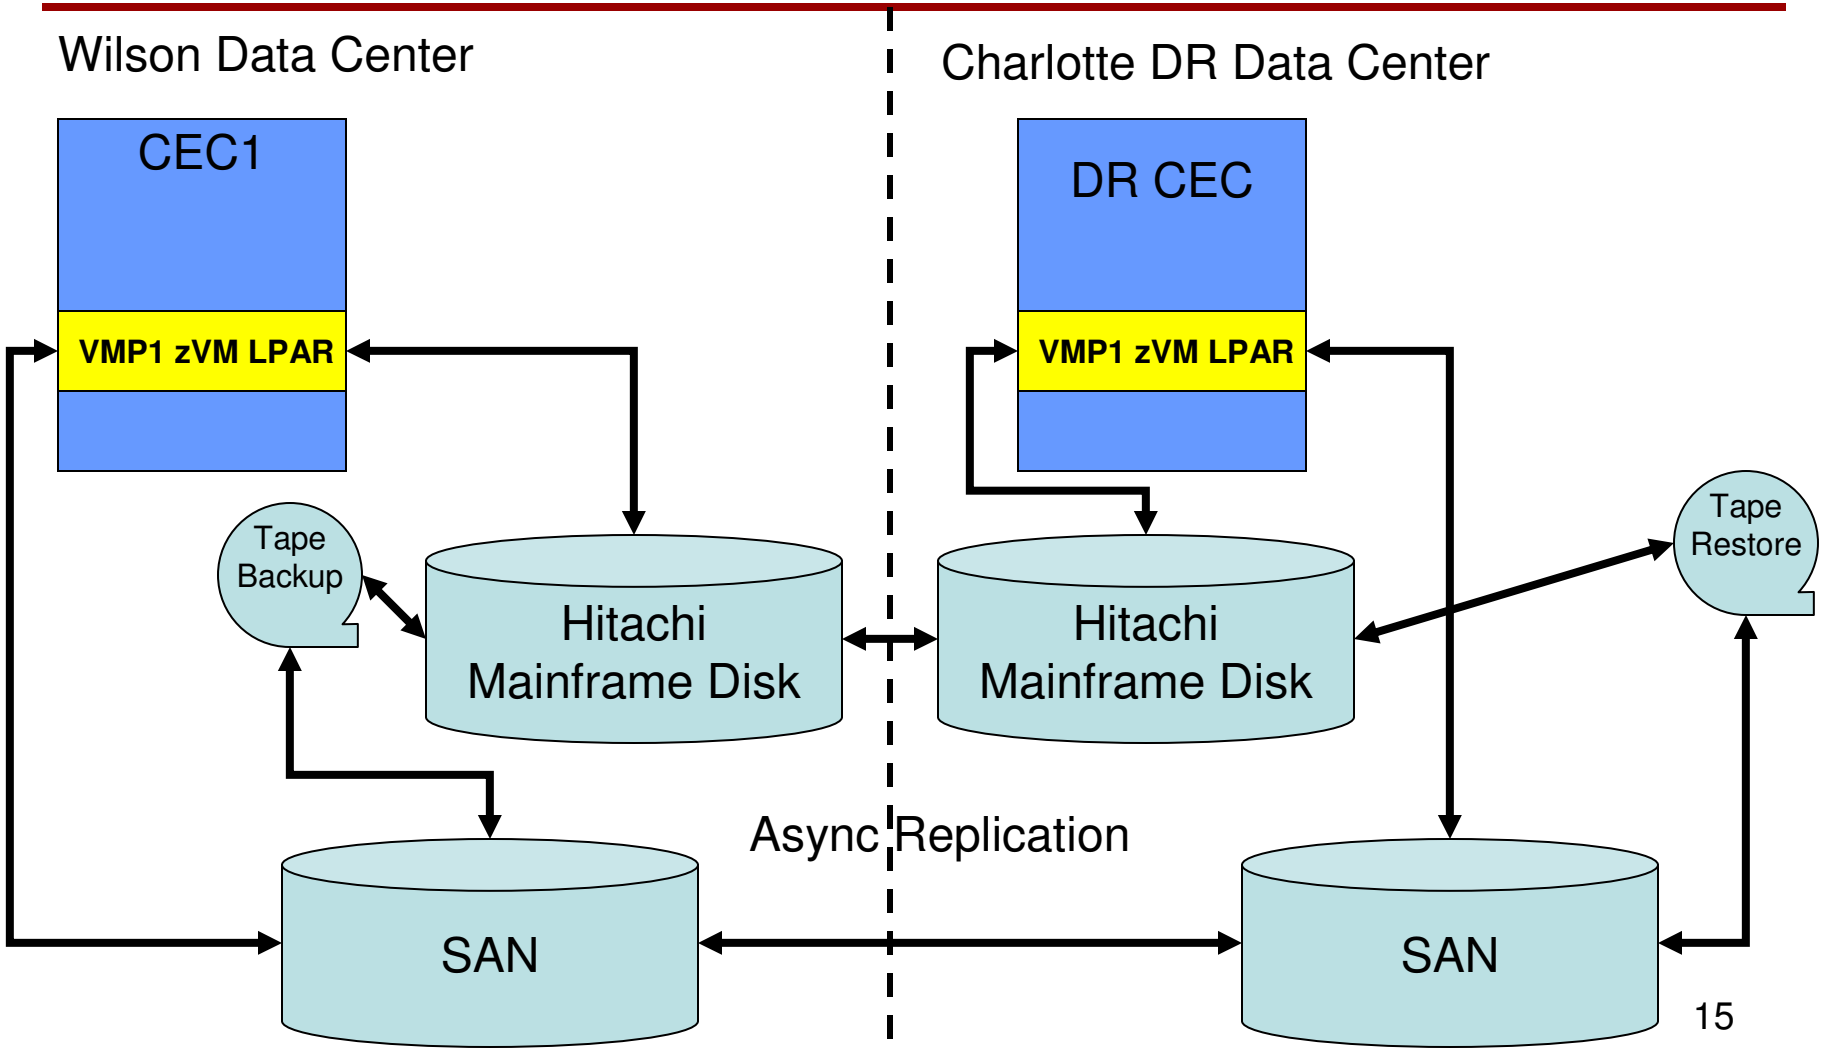

4. Workloads

BB&T Best Bank In Town Since 1872

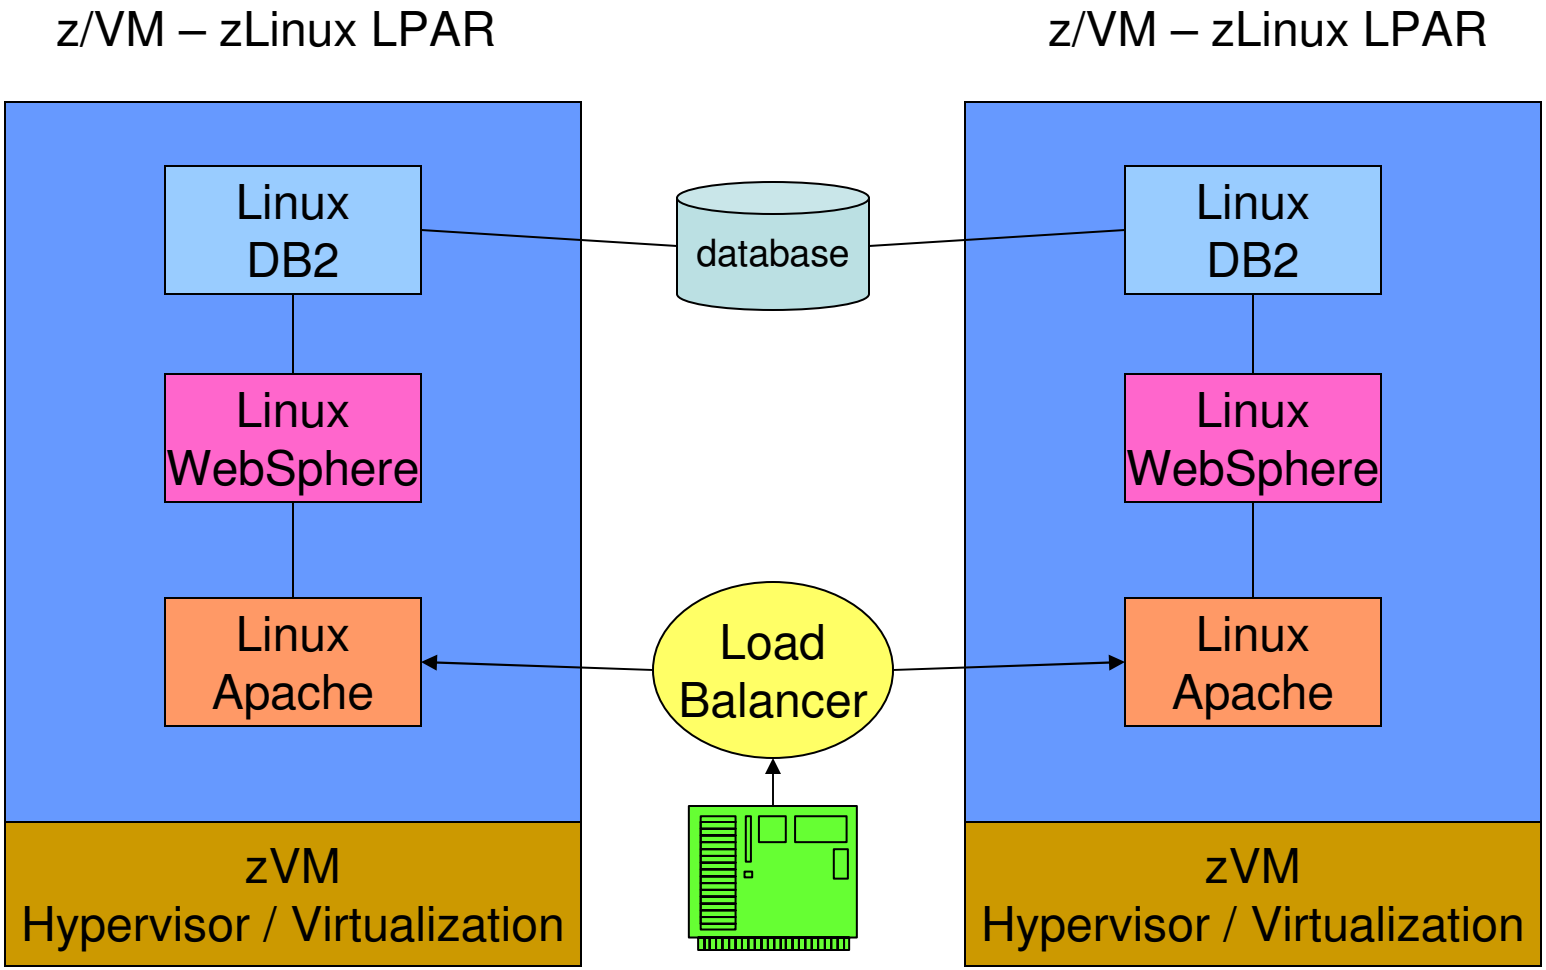

4. Applications

BB&T Best Bank In Town Since 1872

- Apache web servers
- WebSphere running Java
  - 64 bit version uses more CPU time than 32 bit version
  - WAS v7 changes the way it dispatches – be very careful on over committing memory with WAS v7
- MQ – excellent fit for zLinux
- Databases
  - DB2 and/or Oracle
  - Huge software license savings

4.1 Workload characteristics

BB&T Best Bank In Town Since 1872

- BB&T standards – “your mileage may vary”
- Workloads on a distributed server, reasonably new hardware (so reasonably fast processors), that fit these criteria are considered “candidates” for zLinux:
  1. <= 8 distributed “cores”, running less than 50% utilization
  2. <= 8 gig of memory for application servers
  3. <= 16 gig of memory for database servers
  4. No limitations based on I/O demands
  5. At BB&T, we use Red Hat, so the application must be supported on RHEL
  6. Vendor applications – must be supported on zLinux
  7. Drivers:
     1. Proximity to mainframe data
     2. Software license savings

4. Workloads

BB&T Best Bank In Town Since 1872

- In our BB&T zLinux gun sights:

- Pega – Business Process Automation and Workflow tool (software license savings)
- Tibco – Enterprise Service Bus (proximity to mainframe applications)
- Informatica – ETL tool, to be used for integration of merger applications (heavy I/O, large database, priximity to mainframe applications)
- Inventory of old P-Series servers running a varitey of Java applications
- All distributed DB2 and Oracle database

5. Staff considerations

BB&T Best Bank In Town Since 1872

- Your zOS systems programmers can support zVM with just a little bit of training
- The same Sysprogs can build the base zLinux instances, then turn them over to a server Systems Administrator
- SA's do their normal stuff – Linux is Linux

- Monitoring – VERY important on shared platforms

- Capacity Planning – VERY, VERY important on shared platforms

6. Weaknesses of System Z Platform

BB&T Best Bank In Town Since 1872

- Hardware is very expensive
  - Compute intensive applications that need multiple dedicated cores will be VERY expensive to run on System Z on any operating system
  - Memory is much more expensive than RISC or x86 – careful with applications that need lots of memory
- x86 is commodity hardware and very inexpensive to have lots of cores (but be careful of sofware licenses)
- RISC will run heavy compute intensive applications better than others
- System Z is very scalable, but has limits. IFLs are relatively inexpensive, until you have to buy a whole book or a whole CEC
- Economies of scale apply – which means the more stuff you put on System Z, the less expensive (TCO) it will become – and also means that it may be cost prohibitive for running a very small number of applications

7. Where do you go from here?

BB&T Best Bank In Town Since 1872

1. Start with a rough cost model – use what I gave you, plug in numbers for your shop. Compare that with your existing costs.
2. Rough estimate of savings
3. Do a Proof Of Concept
4. While doing the POC, look for candidate workloads
5. Socialize with your application teams – very important to have them on board
6. Start with test applications only – no DR to worry about, low risk, easy
7. Educate your staff
8. Work your way up to bigger and better

BB&T Best Bank In Town Since 1872