

Planning for the Elimination of ESCON I/O on System z

Aug 11, 2011; 8:00-9:00

Session # 9899

Dr. Steve Guendert
Brocade



Patty Driever
IBM



Sean Seitz
Optica



Speakers



Steve Guendert, Ph.D.

Principal Engineer

Brocade Communications



Patty Driever

*System z I/O and
Networking Technologist*

IBM Corporation



Sean Seitz

*Vice President
Technical Services*

Optica Technologies

Agenda

- IBM System z ESCON Roadmap (*Patty Driever*)
- Intro to Managed Evolution for System z (*Sean Seitz*)
 - Current State of ESCON utilization
 - What is Managed Evolution?
- The Managed Evolution solution (*Sean Seitz*)
 - Managed Evolution architecture and topology
 - Managed Evolution applications and customer examples
- FICON Migration and Optimization Tools (*Steve Guendert*)
 - Channel consolidation considerations
 - Tools for use in your migration analysis
- Questions?

System z ESCON Roadmap

Patty Driever
System z I/O and Networking Technologist



ESCON Statements of Direction

- **Released April 28, 2009**
 - It is IBM's intent for ESCON channels to be phased out. System z10 EC and System z10 BC will be the last server to support greater than 240 ESCON channels.
- **Released July 22, 2010**
 - **The IBM zEnterprise 196 will be the last high-end server to offer ordering of ESCON channels.**
- **Released February 15, 2011**
 - **The IBM zEnterprise 196 (z196) will be the last high-end server to support ESCON channels:** IBM plans not to offer ESCON channels as an orderable feature on high-end System z servers which follow the z196 (machine type 2817). In addition, ESCON channels **cannot be carried forward** on an upgrade to such a follow-on server. This plan applies to channel path identifier (CHPID) types CNC, CTC, CVC, and CBY and to feature code numbers 2323 and 2324. **System z customers should continue migrating from ESCON to FICON.** Alternate solutions are available for connectivity to ESCON devices.
- **Released July 12, 2011**
 - **The IBM zEnterprise 196 (z196) and the IBM zEnterprise 114 are the last System z servers to support ESCON channels:** IBM plans not to offer ESCON channels as an orderable feature on System z servers which follow the z196 (machine type 2817) and z114 (machine type 2818). In addition, ESCON channels **cannot be carried forward** on an upgrade to such follow-on servers. This plan applies to channel path identifier (CHPID) types CNC, CTC, CVC, and CBY and to feature code numbers 2323 and 2324. **System z customers should continue to eliminate ESCON channels from the mainframe wherever possible.** Alternate solutions are available for connectivity to ESCON devices.

ESCON Facts

- ESCON cards support ESCON (CNC), ESCON CTC (CTC), Block Multiplexor (CVC) and Byte (CBY) channel types
- The majority of System z installs include some amount of ESCON channels
 - Configured as one of the above channel types
 - *Block Multiplexor and Byte channels and their associated devices are not extinct*
- Q: So what to do? A: Embark on a “Managed Evolution”
 - IBM Global Technology Services offers an ESCON to FICON Migration solution, Offering ID #6948-97D, to help facilitate migration from ESCON to simplify and manage a single physical and operational environment while maximizing green-related savings.

Managed Evolution for System z

- **What is Managed Evolution for System z?**
 - Managed Evolution for System z is a strategic infrastructure simplification solution that aligns with IBM's **New Enterprise Data Center** and **Green Leadership** initiatives.
- **Managed Evolution *conditions the System z environment* enabling customers to:**
 - Simplify the cutover to a new z platform while becoming “new workload ready”
 - **System z Recommended Best Practice:** Deploy 100% FICON channels on the host and exploit the full benefits of FICON:
 - Maintain access to ESCON and Bus and Tag devices required to support key applications
 - Consolidate ESCON infrastructure and operations
 - IBM Global Technology Services offers an ESCON to FICON Migration solution, Offering ID #6948-97D, to help facilitate migration from ESCON to simplify and manage a single physical and operational environment while maximizing green-related savings.
- **IBM System z and Optica Technologies have collaborated to deliver the exclusive technology (Prizm) required to support this strategy**
 - Prizm completed a System z qualification for all ESCON devices in July 2009
 - Prizm works direct attached to a System z or with any System z qualified director

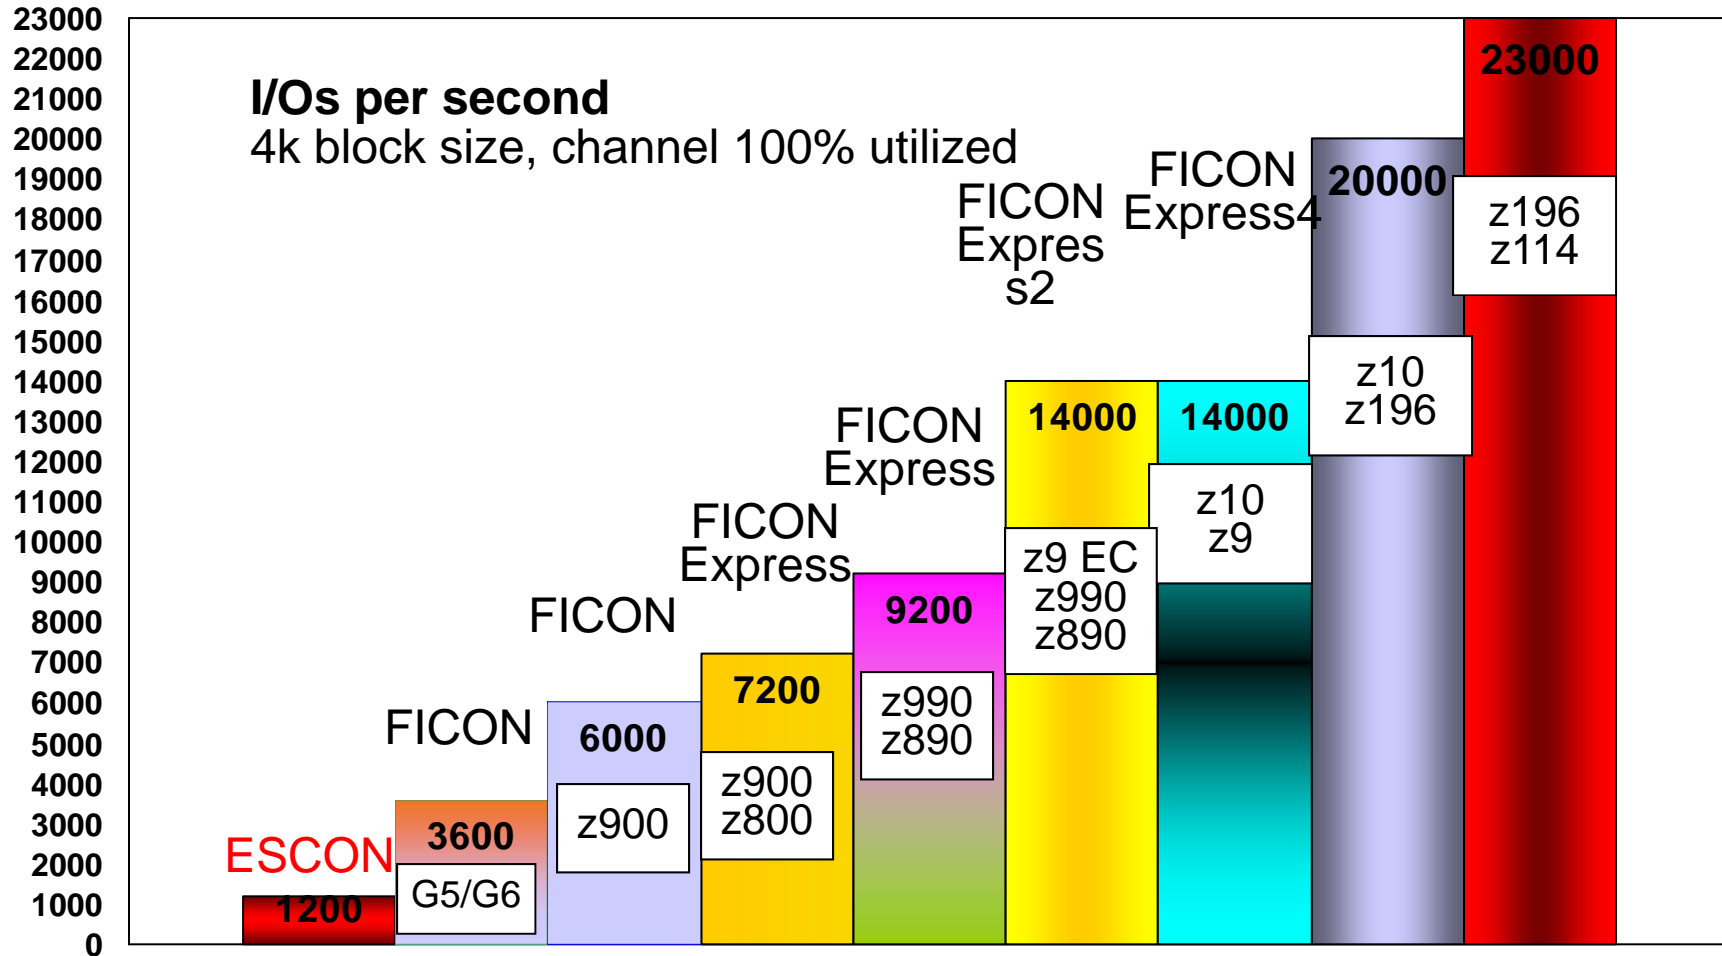
Strategic Benefits of FICON

- Improved workload management
 - Measurements – granular to service class
 - Allows algorithms for WLM based I/O priority, DCM & intelligent data placement
- I/O start rate and bandwidth performance
 - Reduction in interlocks between channel and control unit
 - Multiplexing of I/O operations (to multiple devices)
 - Pipelining of I/O operations (to a single device)
 - Improved (less chatty) protocol on the link with zHPF
 - *Initially released on z10 systems*
 - *Supported by major enterprise-class DASD control unit vendors*
 - I/O priority
 - MIDAWs/TIDAWs
- Extended distance
 - ESCON supports maximum of 9km distances without data droop
 - FICON supports 100km
 - Enhancements for XRC acceleration
- Multiplexing of mixed workloads
 - ESCON (with Prizm), FICON, zHPF



FICON performance – Start I/Os

Historical Actuals

FICON Express8 FICON Express8S



FICON performance on System z



Introduction to Managed Evolution

Sean Seitz
VP Technical Services



The Current State of ESCON Utilization



The Reality Is.....


- 85% of all System z customers and 78% of all System z machines in use today have active CNC (native ESCON) channels defined.
- 33% of all z9 and z10 systems have CVC and CBY (converted mode ESCON) channels defined to support Bus/Tag devices.
- Vast majority of large mainframe customers still rely on ESCON and/or Bus/Tag based applications.

Sources: zJournal "Mainframe Evolution Survey"
and IBM System z Brand

zJournal Survey: Stable, Legacy Applications Remain on ESCON



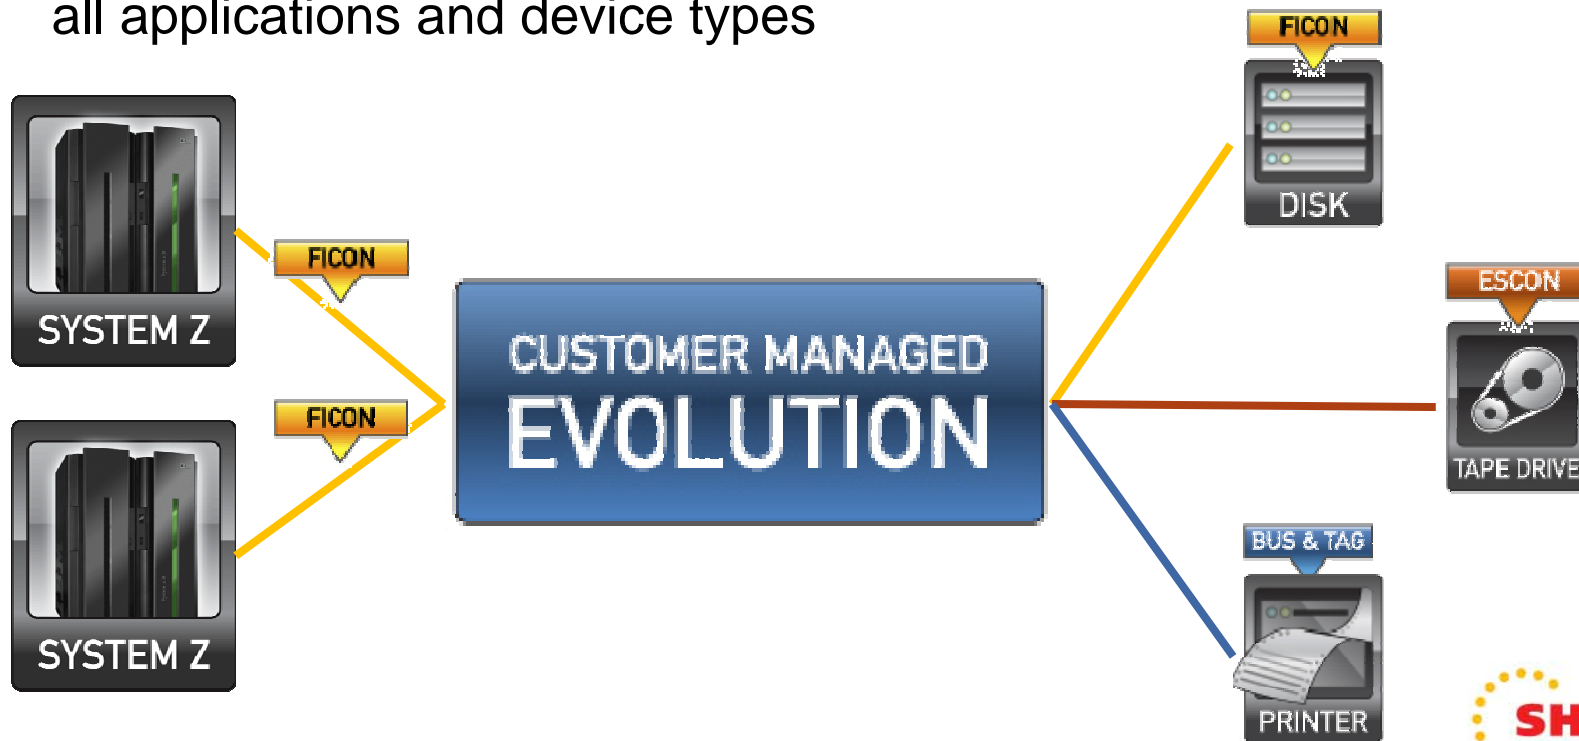
“What factors compel you to preserve investments in ESCON applications and device types?”



Source: zJournal “Mainframe Evolution Survey”

What is Managed Evolution for System z?

- Strategically invest in System z / FICON host infrastructure modernization
- Manage your storage and other device types based on application characteristics (FICON, ESCON, Parallel)
- Migrate to a simplified host-based FICON infrastructure supporting all applications and device types



Prizm is the building block...

- Prizm is a purpose built appliance designed **exclusively** for IBM System z
- Prizm converts native FICON (FC) protocol to native ESCON (CNC) protocol allowing ESCON and B/T devices to connect to FICON channels

Front View




PRIZM™

FICON Ports:
LC Duplex

ESCON Ports:
MT-RJ

2U



Rear View

Prizm Basics...

- Prizm is a 2u rack mountable system which converts 1 or 2 FICON channels into 4, 8 or 12 ESCON channels.
 - Prizm also supports bus/tag device attachment via ESBT module
- Prizm is available in the following configurations:
 - 1 FICON (IN) to 4 ESCON (OUT) = 1:4
 - 2 FICON (IN) to 8 ESCON (OUT) = 2:8
 - 2 FICON (IN) to 12 ESCON (OUT) = 2:12
 - Available with long-wave (LX) or short-wave (SX) FICON optics
 - “ESBT” module for support of parallel (bus/tag) devices
- Prizm is easy to configure and install and will attach to a broad array of ESCON (and Bus / Tag) devices.
 - Qualified in the IBM Vendor Solutions Lab in Poughkeepsie, NY
 - Brocade and Optica perform continual interop testing in our test labs

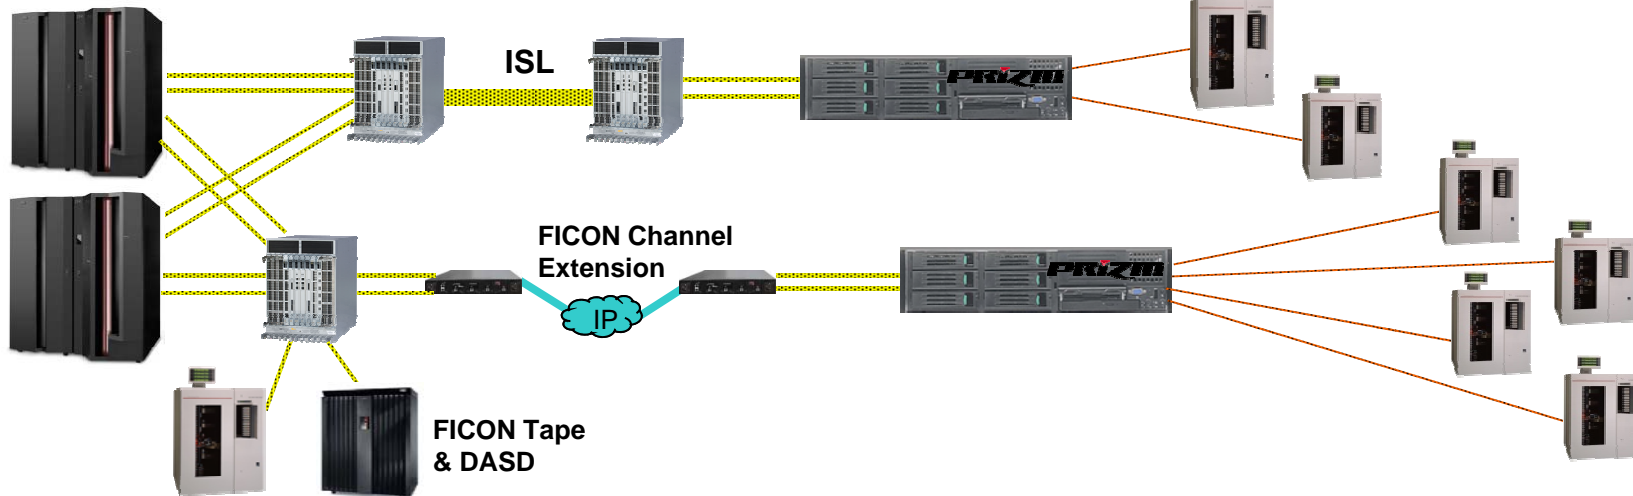
Where does Prizm fit in the infrastructure?

1. Point to Point



2. Switched

3. Cascaded and Channel Extended

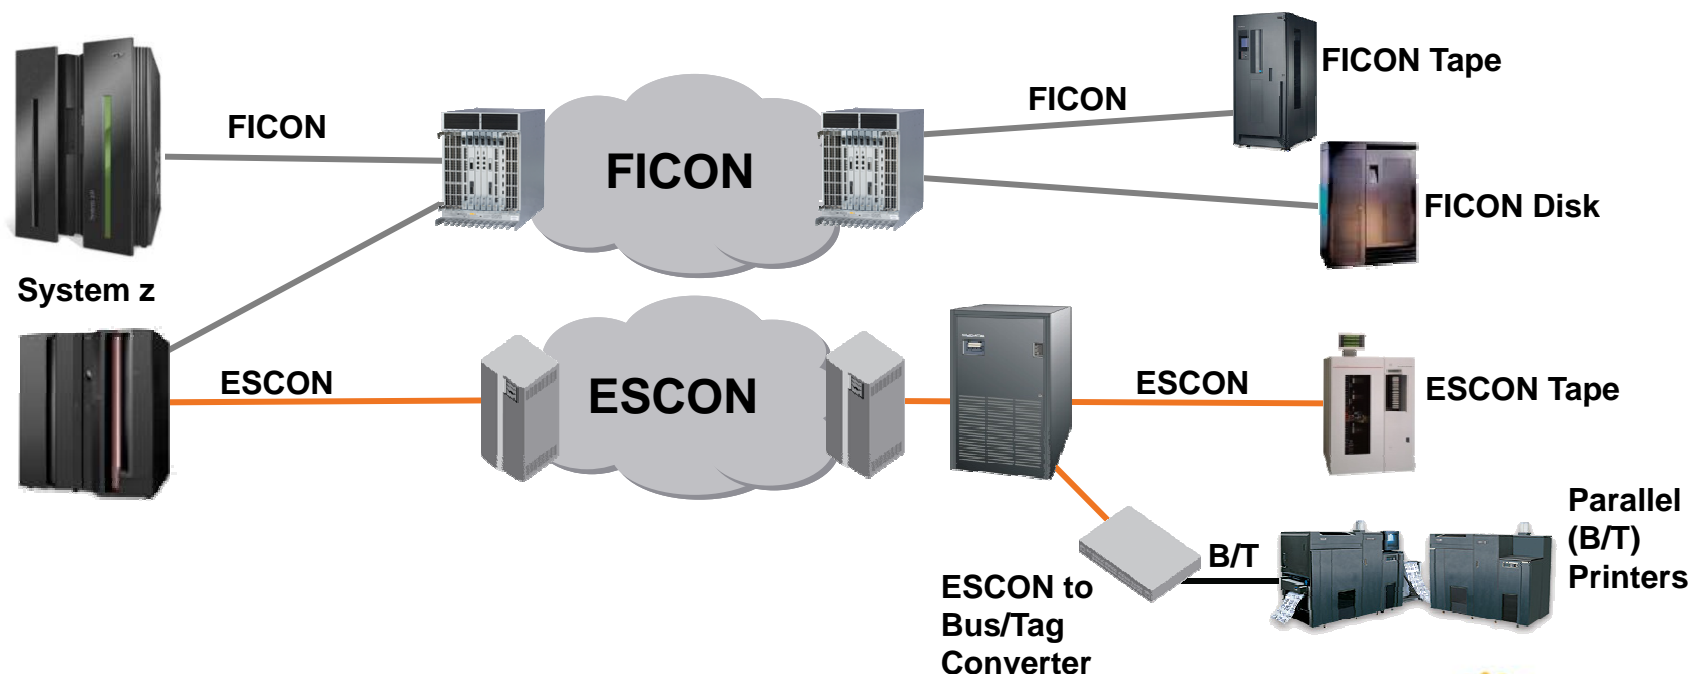


4. Support for a broad set of ESCON and B/T devices: Tape, Printers, Com Devices, FEPs etc.

Today: Two Infrastructures

Current environment

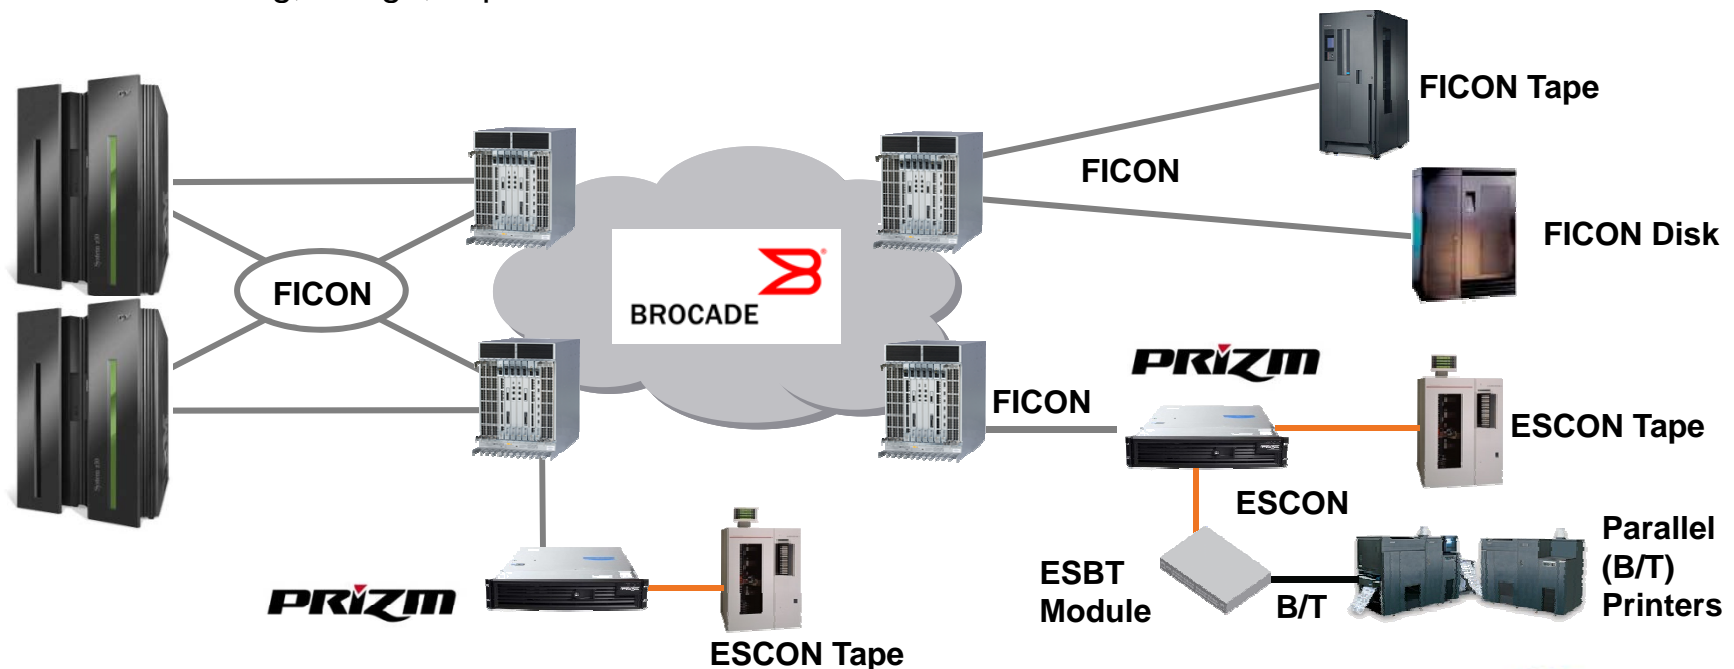
- Dual infrastructures (FICON and ESCON)
- Local and extended distance (ESCON)
- FICON Disk and Tape
- ESCON Tape/Controllers
- Parallel Printers/Controllers



Managed Evolution Architecture

Unified, Thoughtfully Planned Migration Management

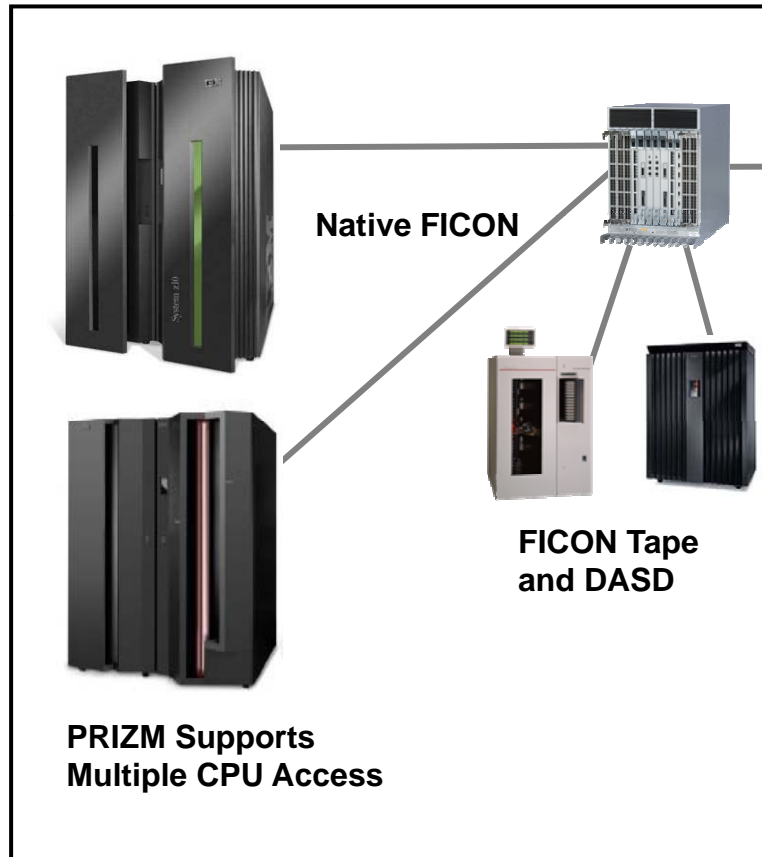
- Single Brocade DCX infrastructure
- Optica PRIZM FICON to ESCON/parallel solutions
- Turnkey Managed Evolution migration management services
 - Planning, design, implementation and maintenance



Managed Evolution Applications

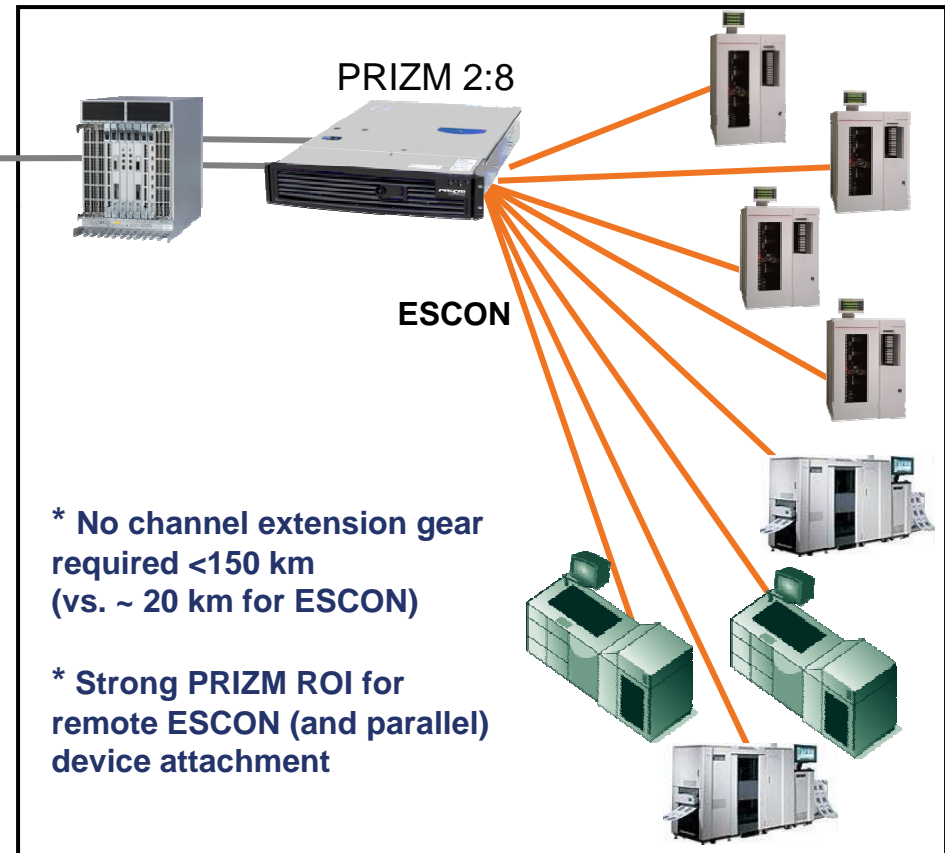
ESCON Device Extension via FICON Infrastructure

Local Site



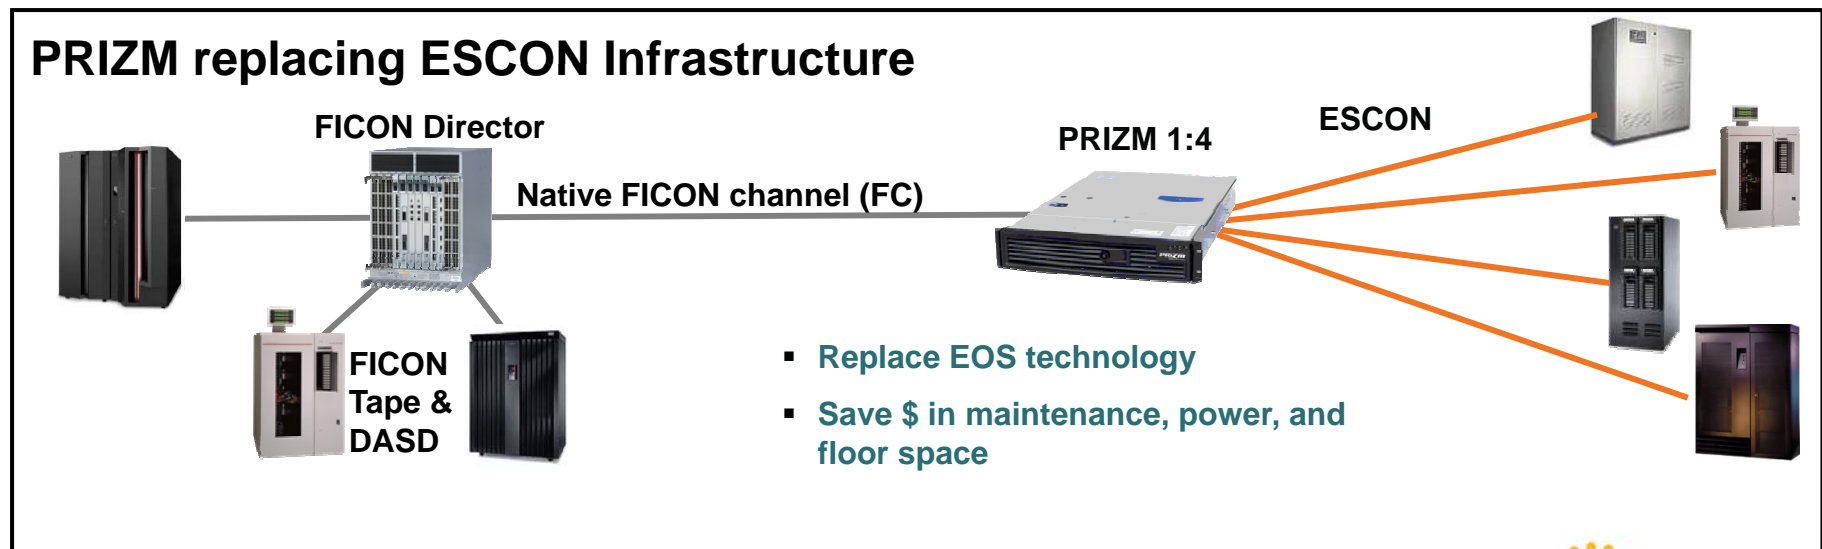

FICON (ISL)

Remote Site



Managed Evolution Applications

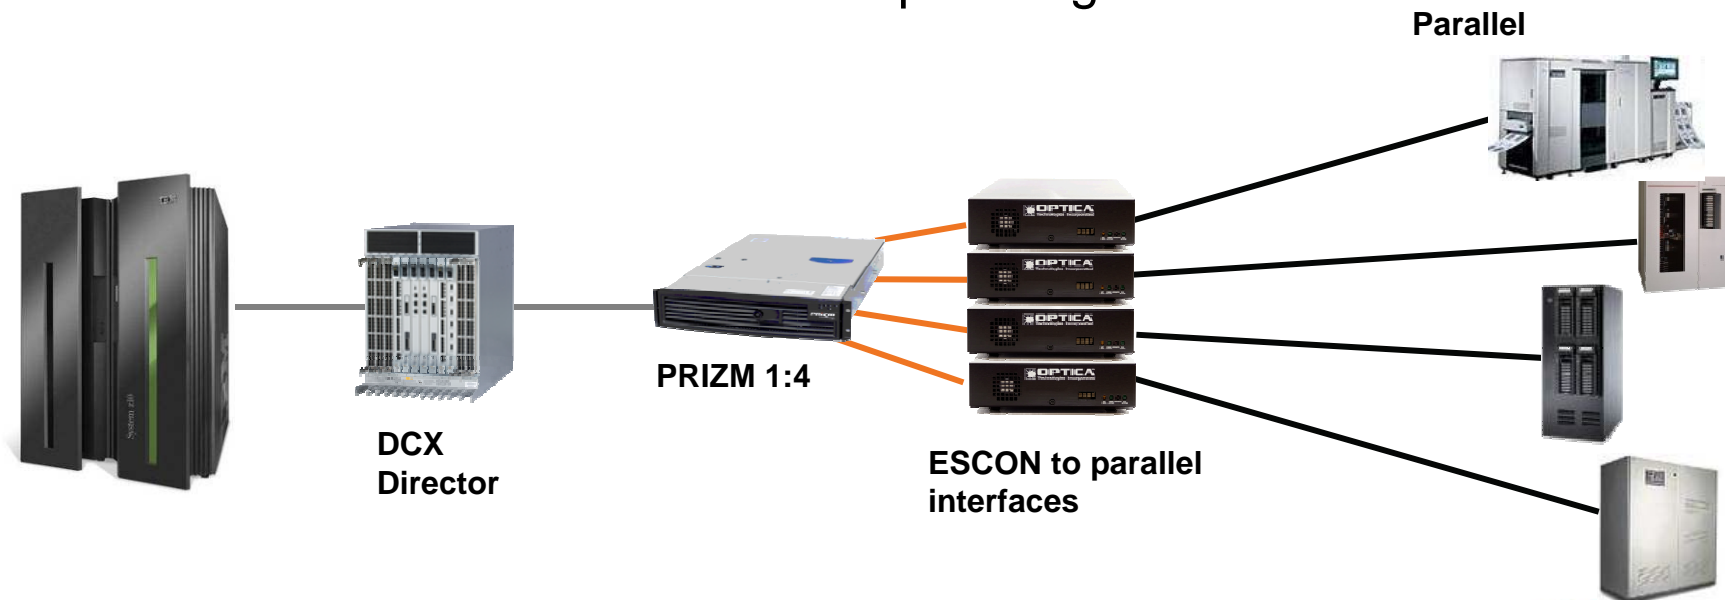
ESCON Director and FICON Bridge Replacement



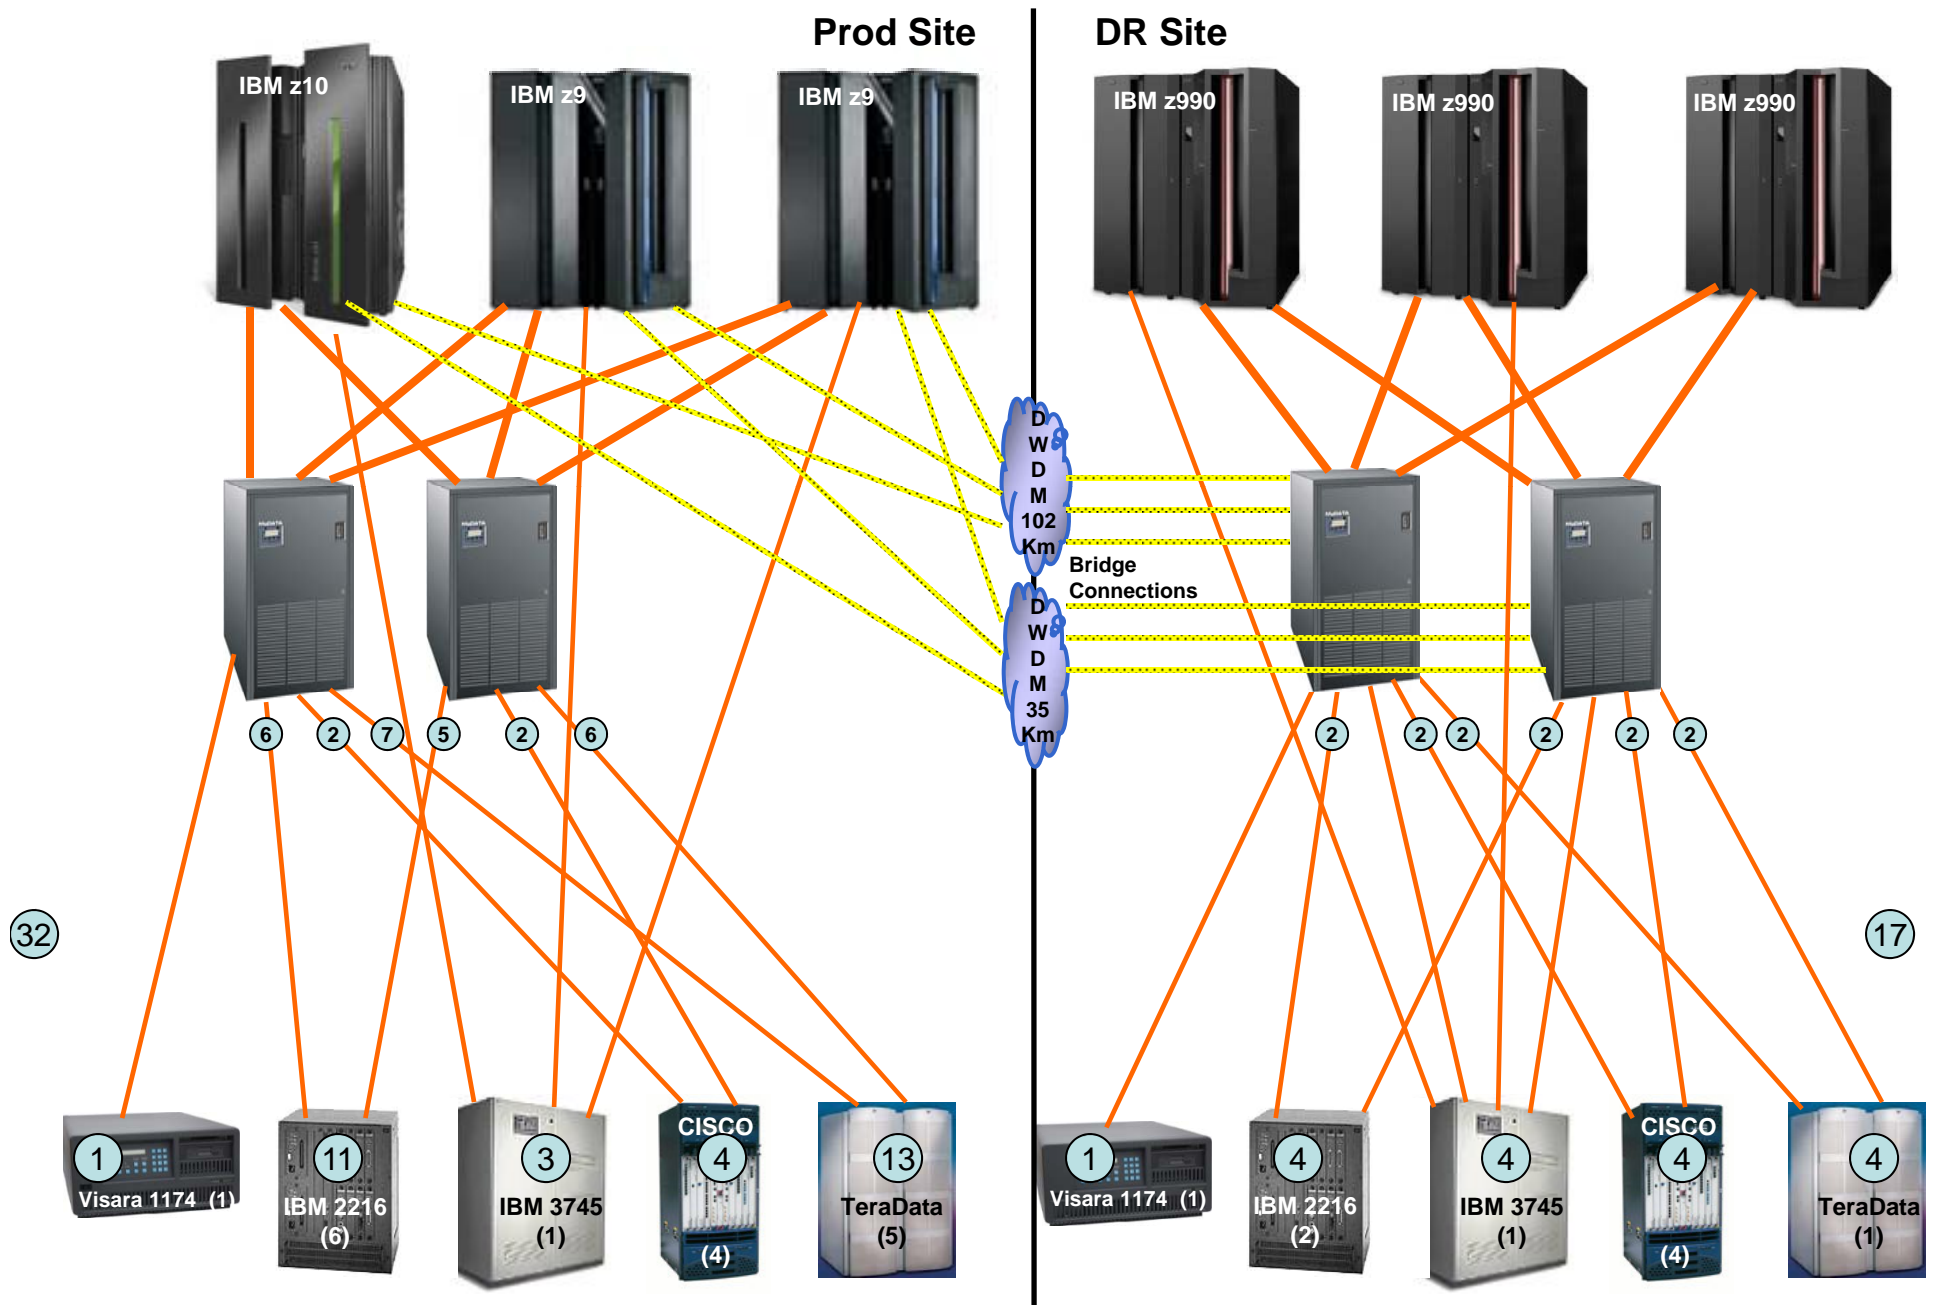
Managed Evolution Applications

Parallel device attachment to FICON channels

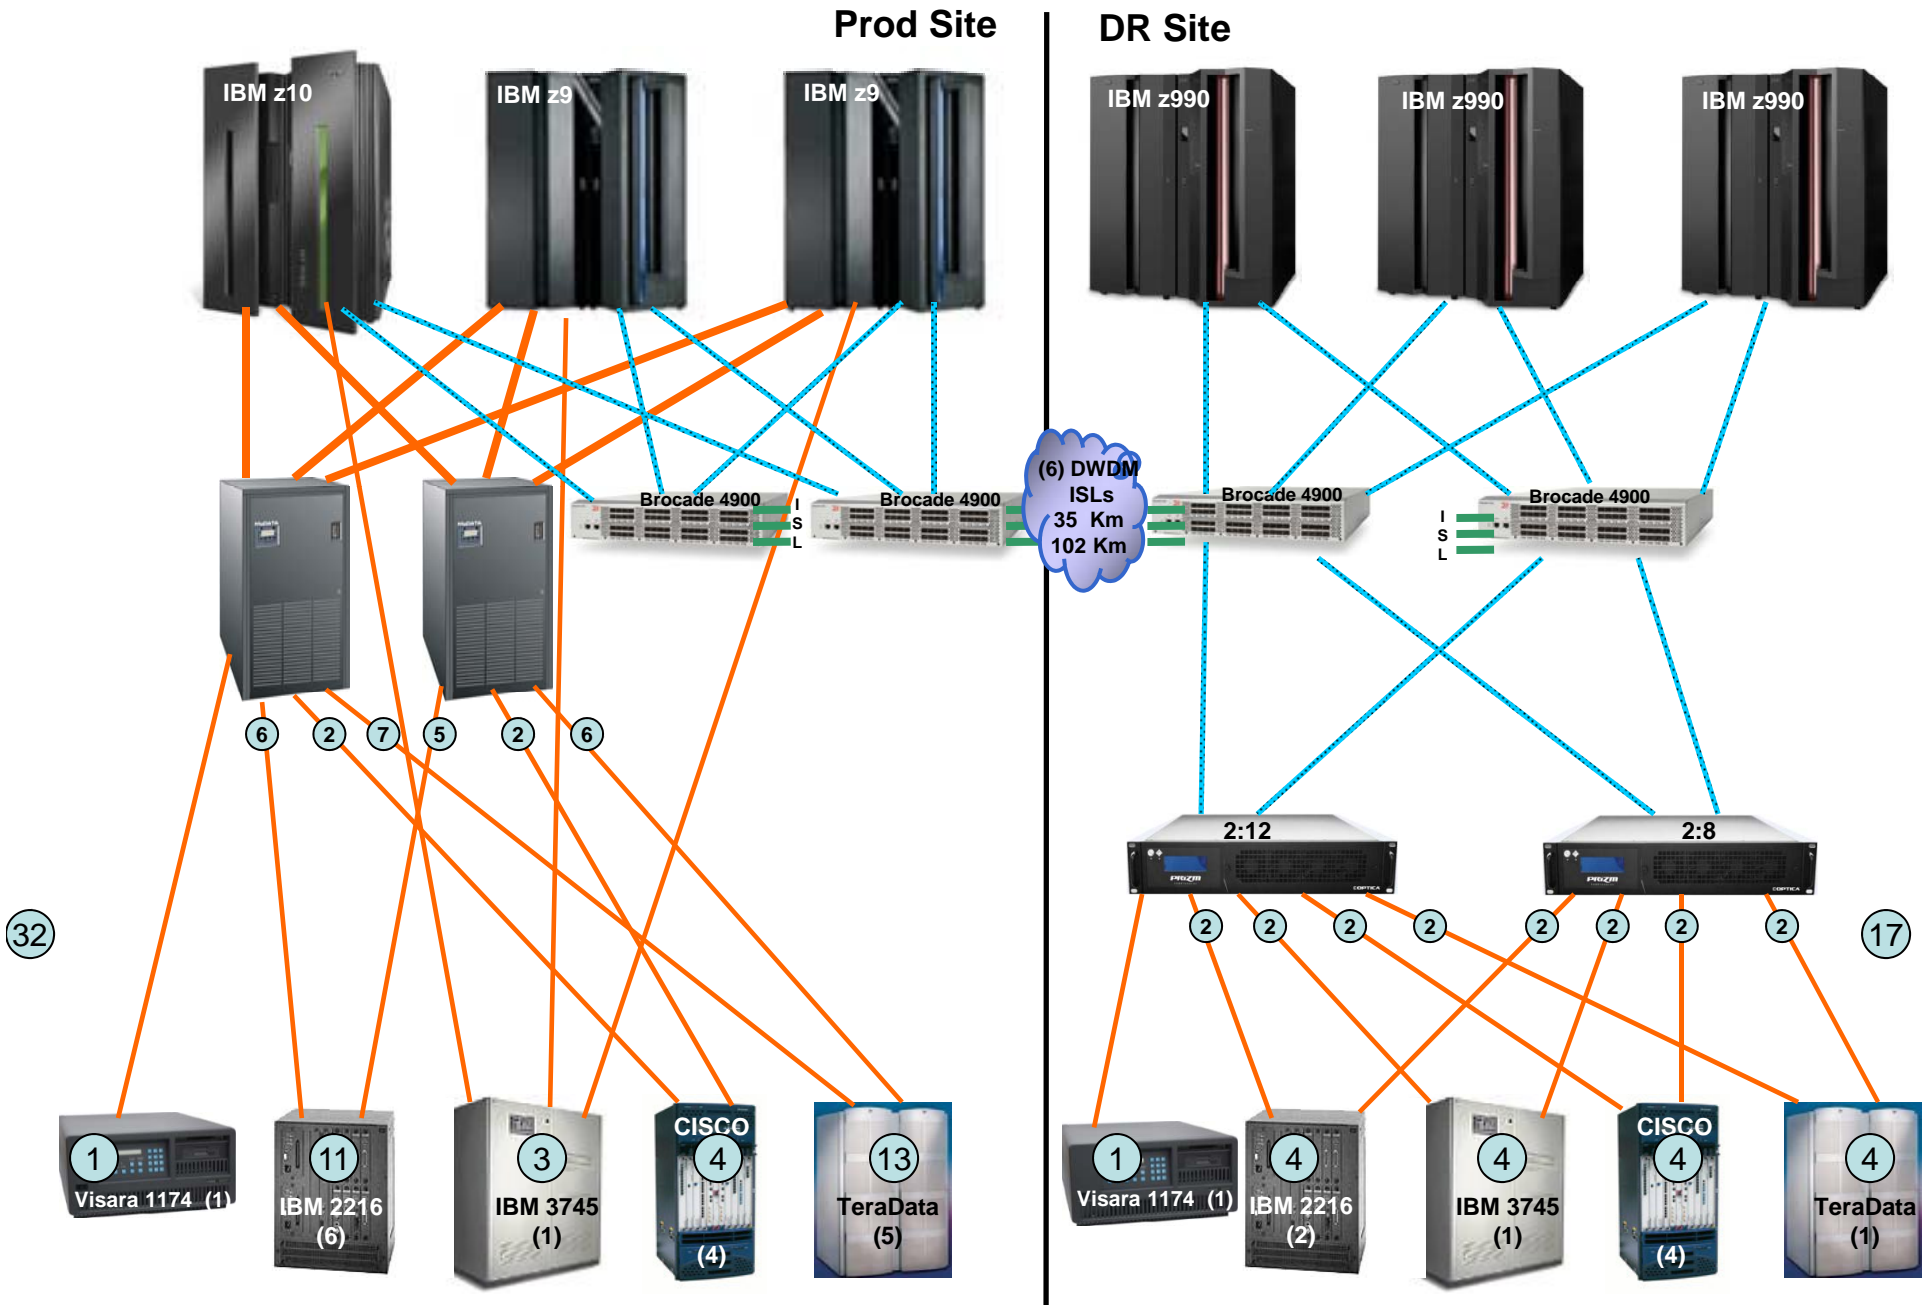
- Maintain installed parallel devices and applications while migrating to FICON on the host
- Parallel device channel extension via FICON
 - Attach parallel devices in remote data centers using FICON InterSwitch Links (ISLs)
- Provides infrastructure options and flexibility for mainframe refresh and new site planning



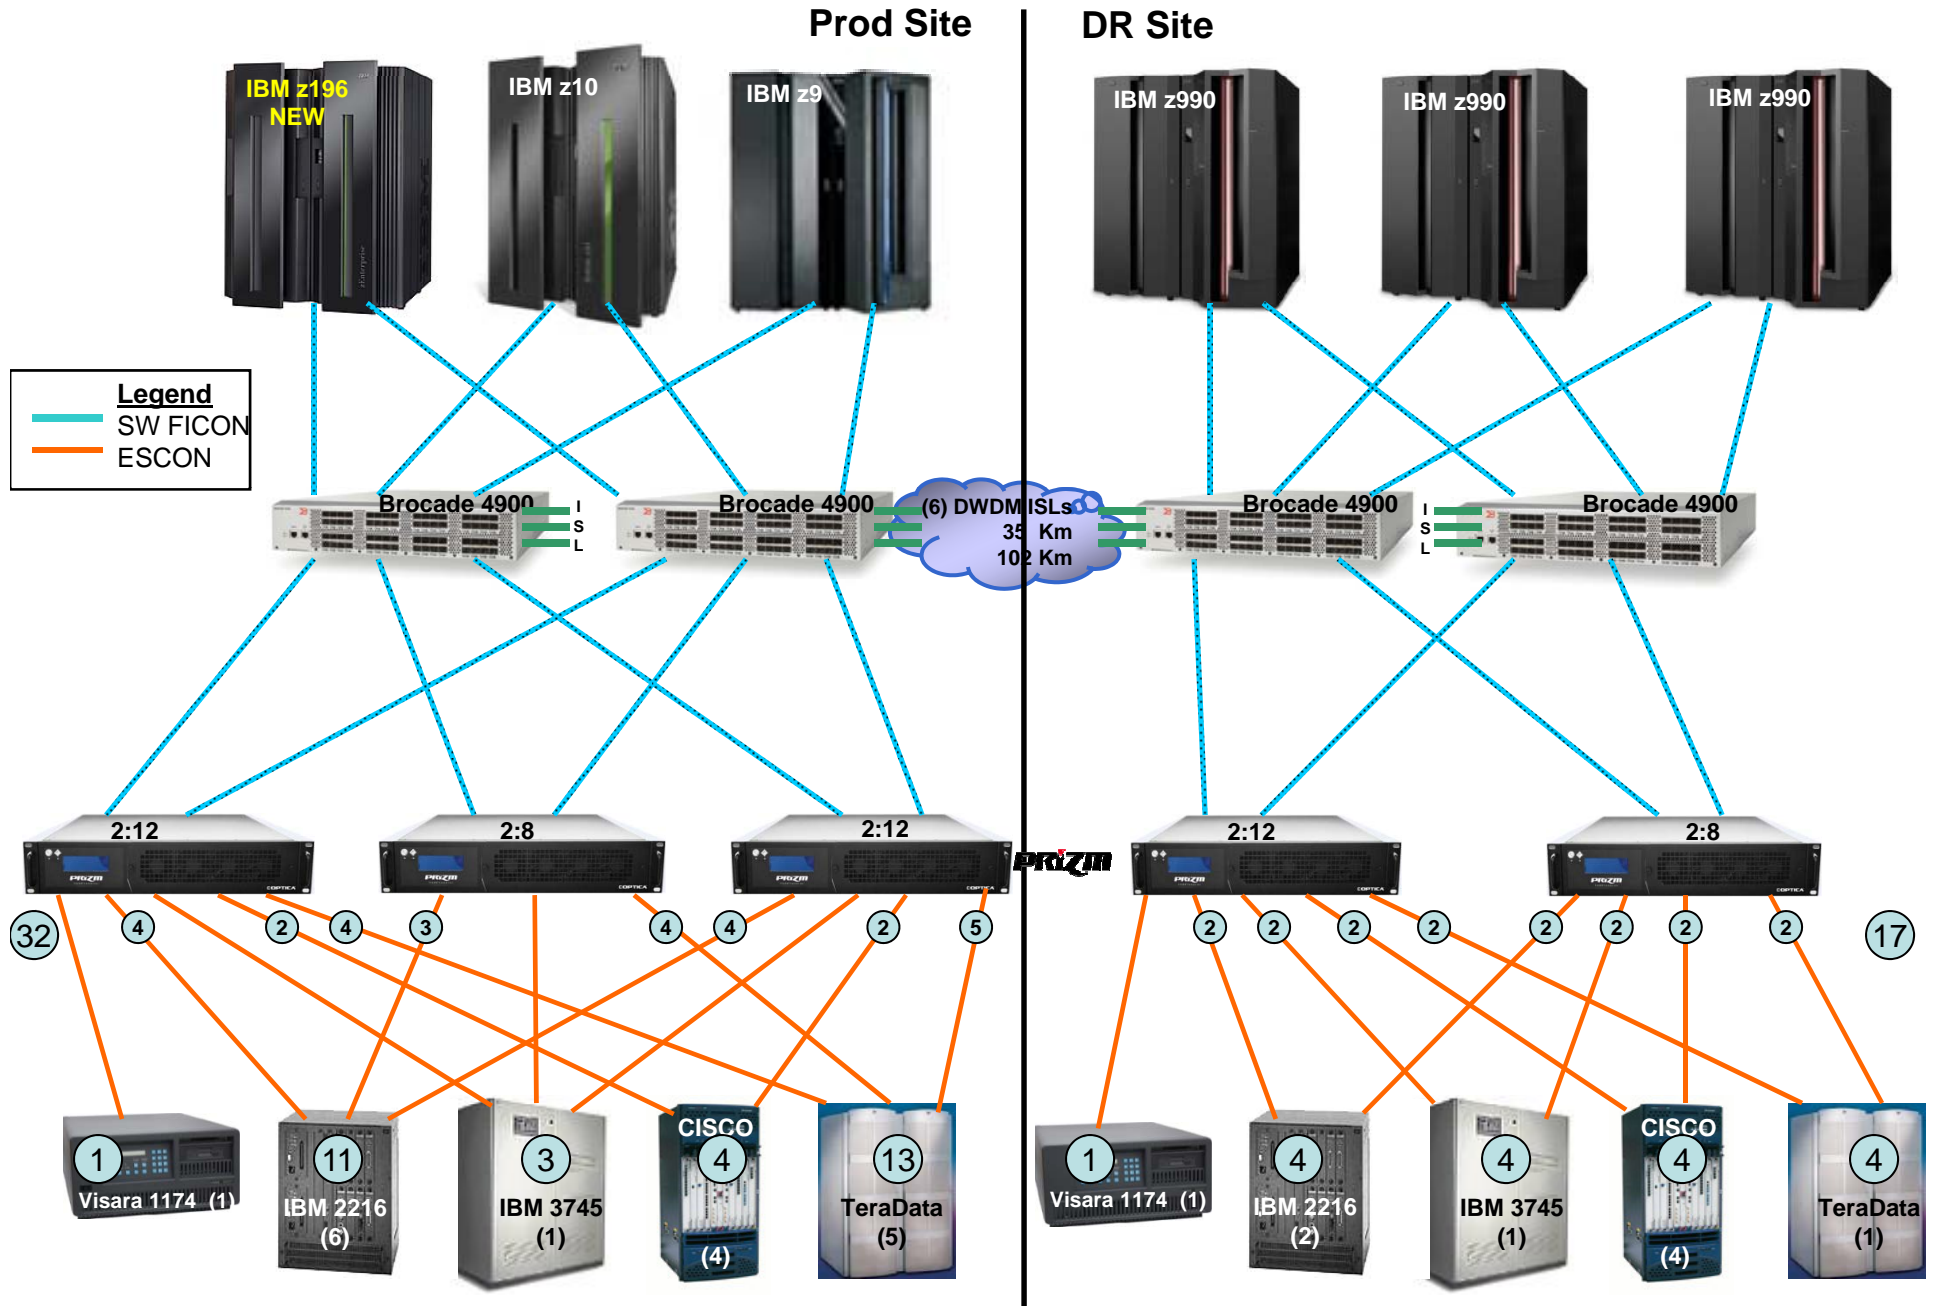
Customer Case Study #1 Original Configuration




Customer Case Study #1 **PRIZM** Phase 1 Design



Customer Case Study #1 **PRIZM** Phase 2 Design (w/ z196)

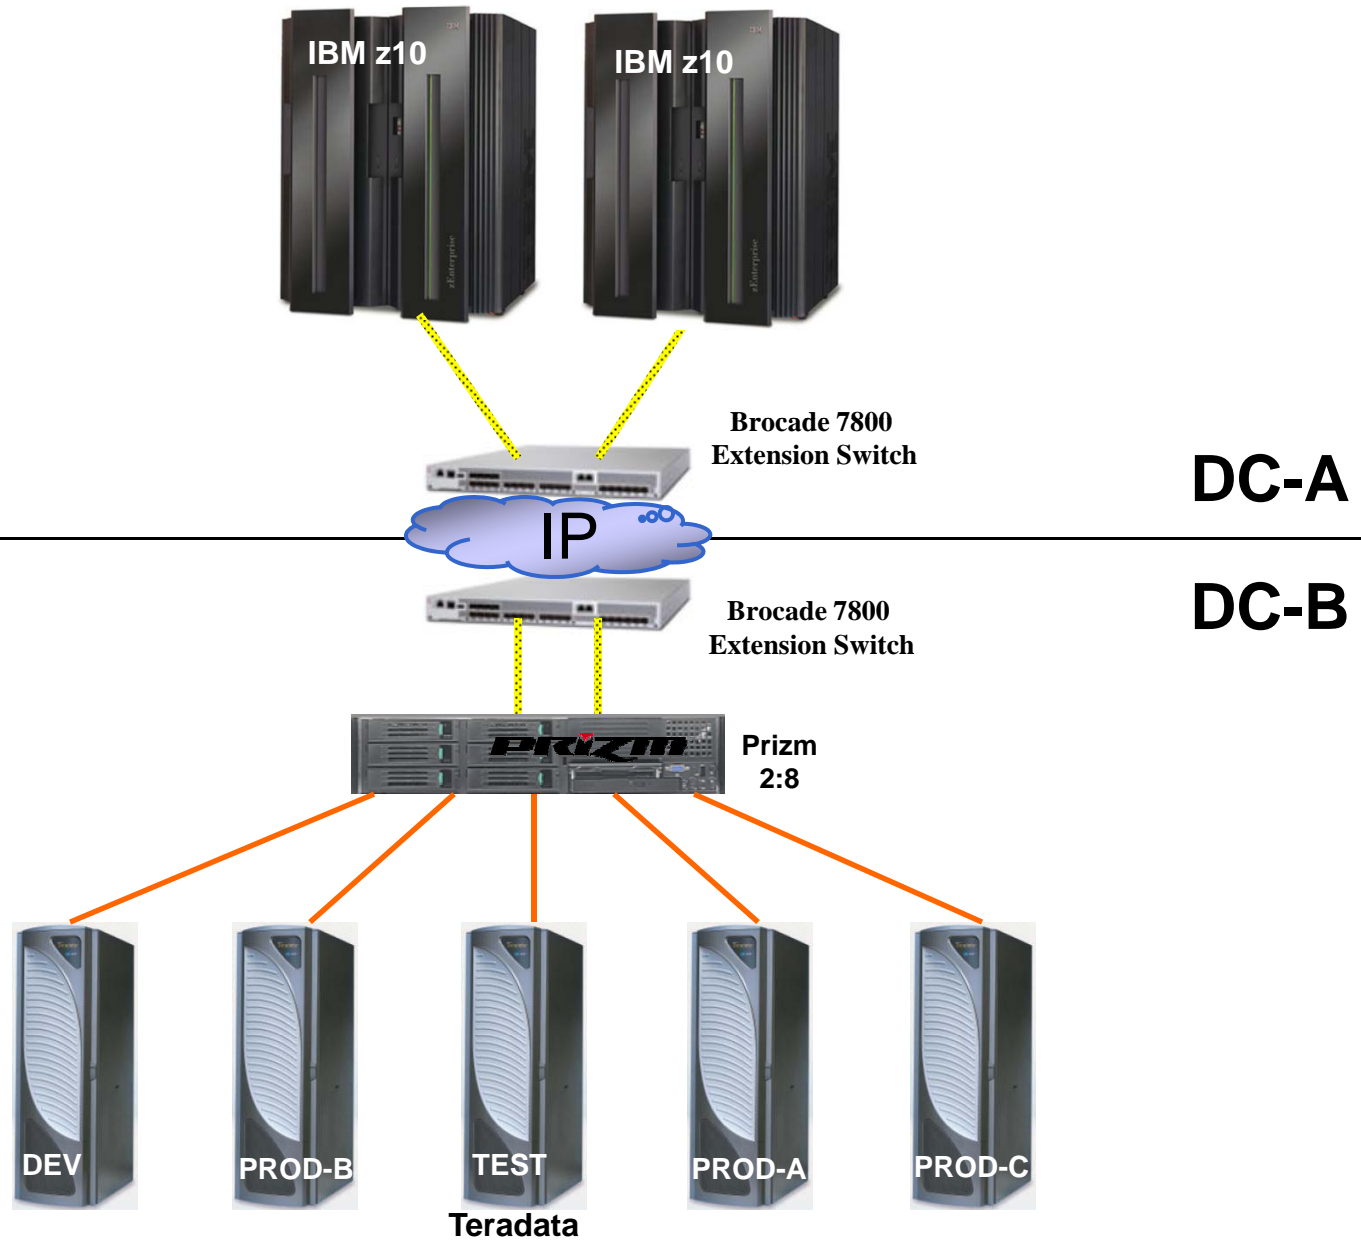


Customer Case Study #2 Original Configuration



Customer Case Study #2

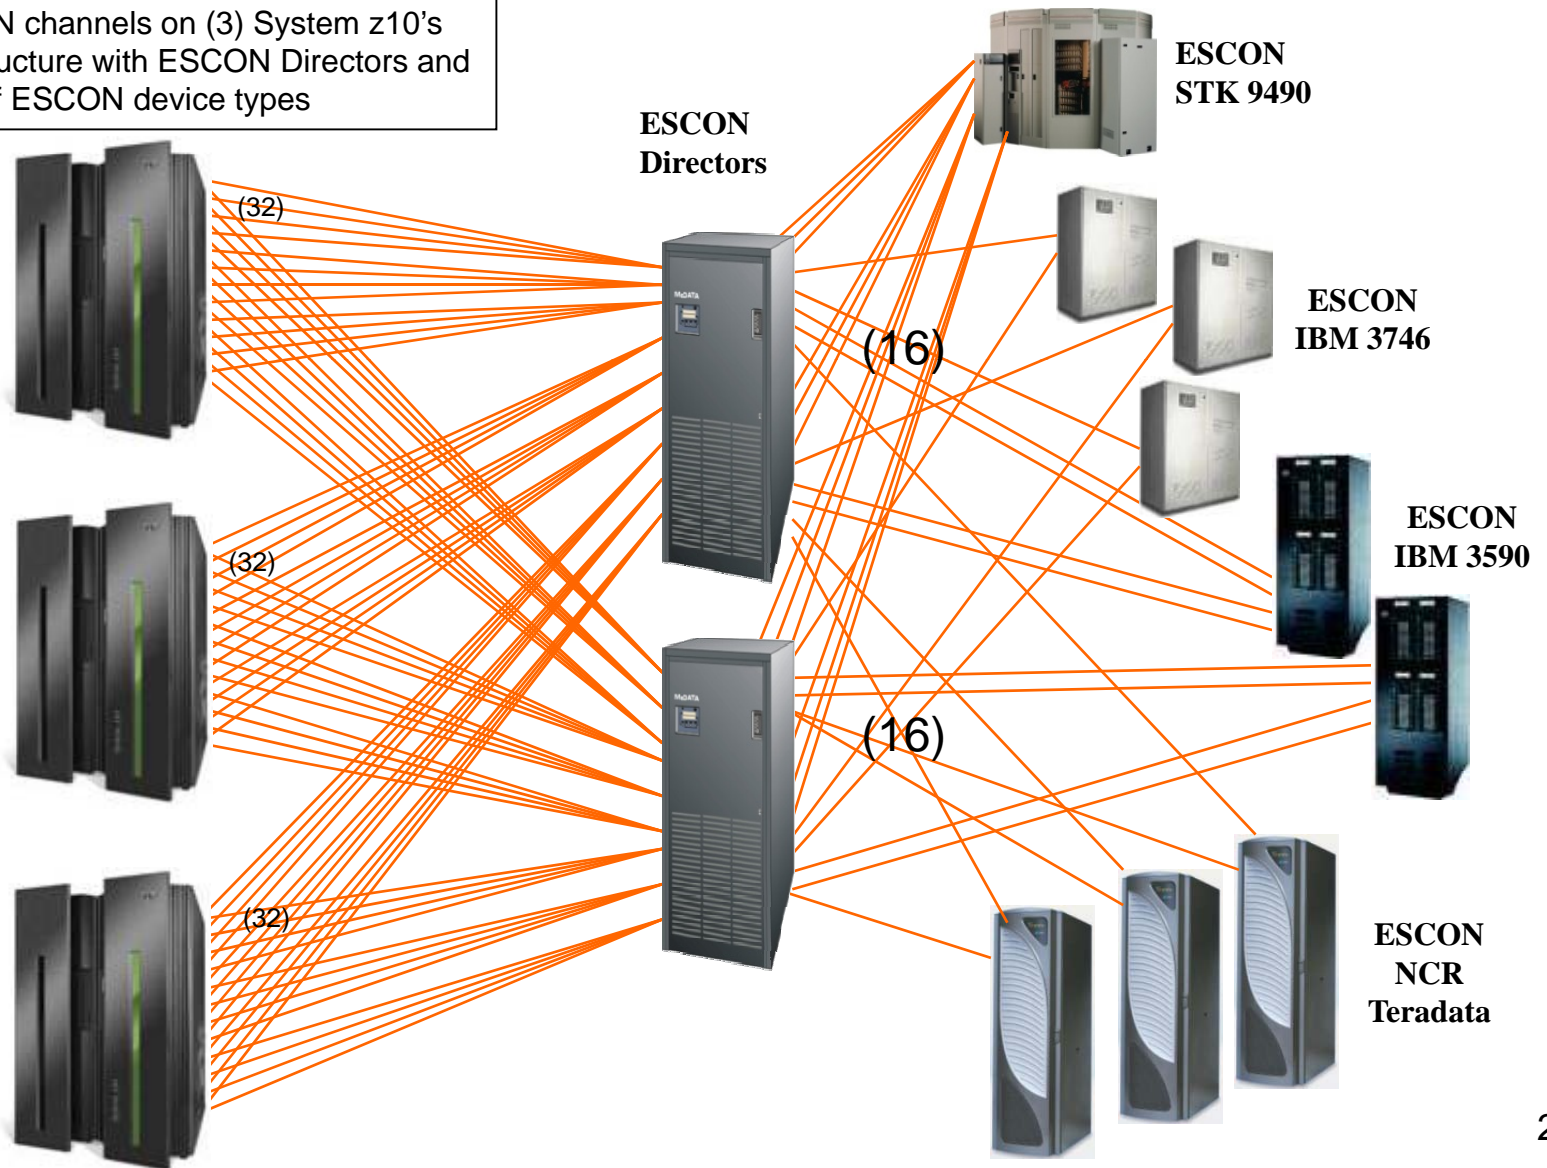
PRIZM Design



Customer Case Study #3 Original Configuration

Original ESCON Configuration

- Host: 96 ESCON channels on (3) System z10's
- ESCON infrastructure with ESCON Directors and a wide variety of ESCON device types

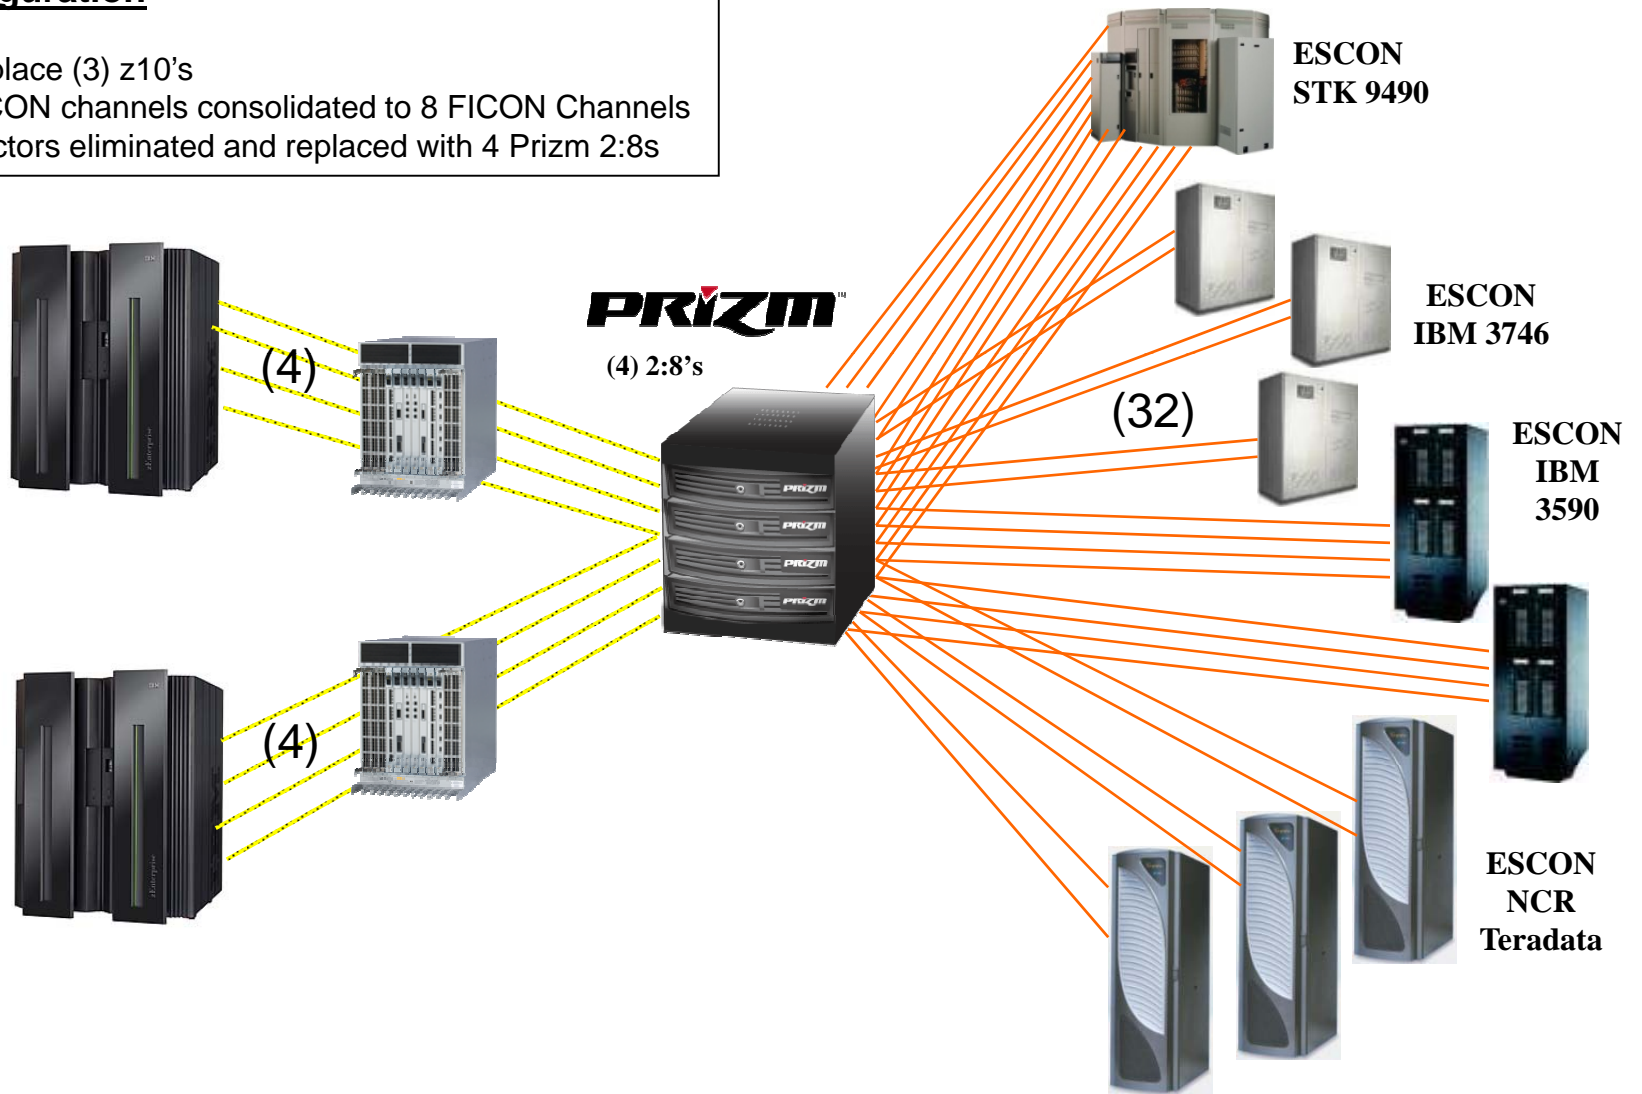


Customer Case Study #3

PRIZM Design

Prizm Configuration

- (2) z196's replace (3) z10's
- Host: 96 ESCON channels consolidated to 8 FICON Channels
- ESCON Directors eliminated and replaced with 4 Prizm 2:8's



Consistent Customer Benefits

- Maximizes value of System z consolidation while reducing the “cutover” risk
- Leverages the value of FICON
- Simplifies I/O and Operations
- Eliminates ESCON as a planning consideration for System z
- Savings on ESCON director maintenance, power, cooling and floor space supports the case for transition

Predominant ESCON Applications for Prizm...

- **Tape Backup/Tape Exchange**
 - Tape has major operational infrastructure and change is expensive
- **SNA Networks/VTAM**
 - Either 37XX or CIP Networks with older ATM Technologies
- **Database Machines**
 - Teradata
- **Print**
 - Print can be ESCON or B/T – major infrastructure – collation, bursting, stacking

NONE of these applications require greater performance

What Events Drive the Change?



- **System z – z196 and z114 Planning and Upgrades**
 - 85% of Mainframe customers have ESCON or a mix of ESCON and FICON today*
 - System z recommends customers plan/move now
- **ESCON Director - Replacement**
 - 1000's of ESCON directors are still in production**
 - Plan for end of service
- **ESCON Channel Extension – Replacement**
 - Over 12,000 nodes deployed**
 - End of Life and Support is here
 - Prizm allows customer to leverage the value of FICON while reducing the cost and complexity of managing ESCON over long distances

Sources:


*System z brand

**Optica estimates

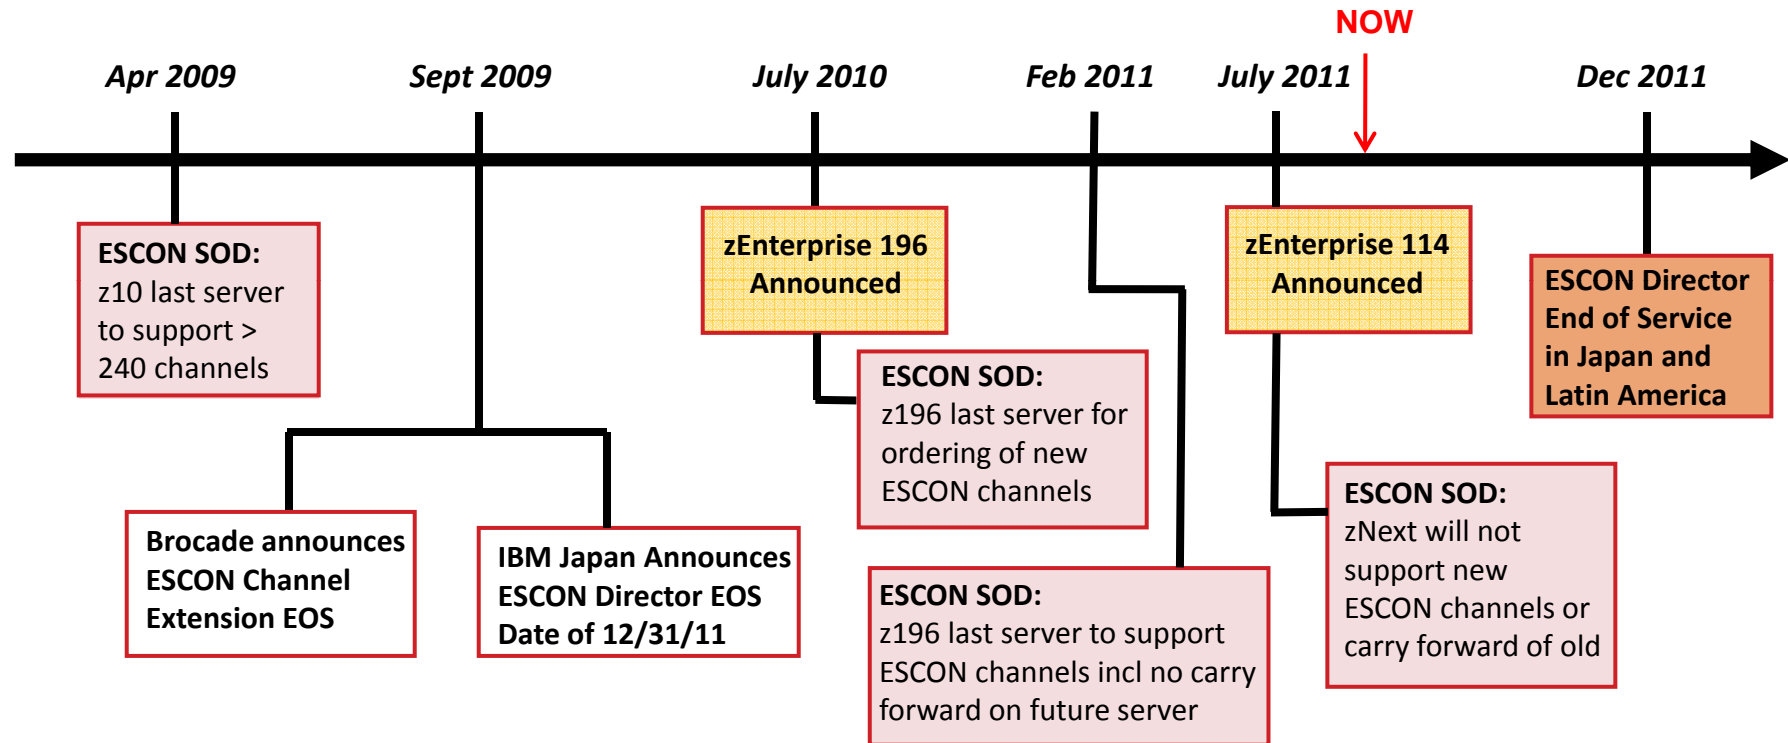
Planning for Prizm (ESCON) Bandwidth is easy!



- System z customers provision I/O with high levels of resiliency
 - 50% I/O headroom or more is the norm to deliver consistent application performance
- Prizm allows customers to share FICON CHIPIDs to service ESCON device requirements and uses a small percentage of available bandwidth
- FICON Express 8 enables customers to eliminate ESCON and consolidate FICON while increasing I/O headroom



The System z ESCON roadmap is clearly defined... - now it's time to make sure you are prepared



Planning for your future System z I/O connectivity....

- Assess current device/application portfolio and requirements
- Assess ESCON switching requirements in preparation for ESCD EOS
- Assess oppty for 8 Gb FICON consolidation
- Consider distance support requirements for ESCON/BT devices



FICON migration and optimization tools

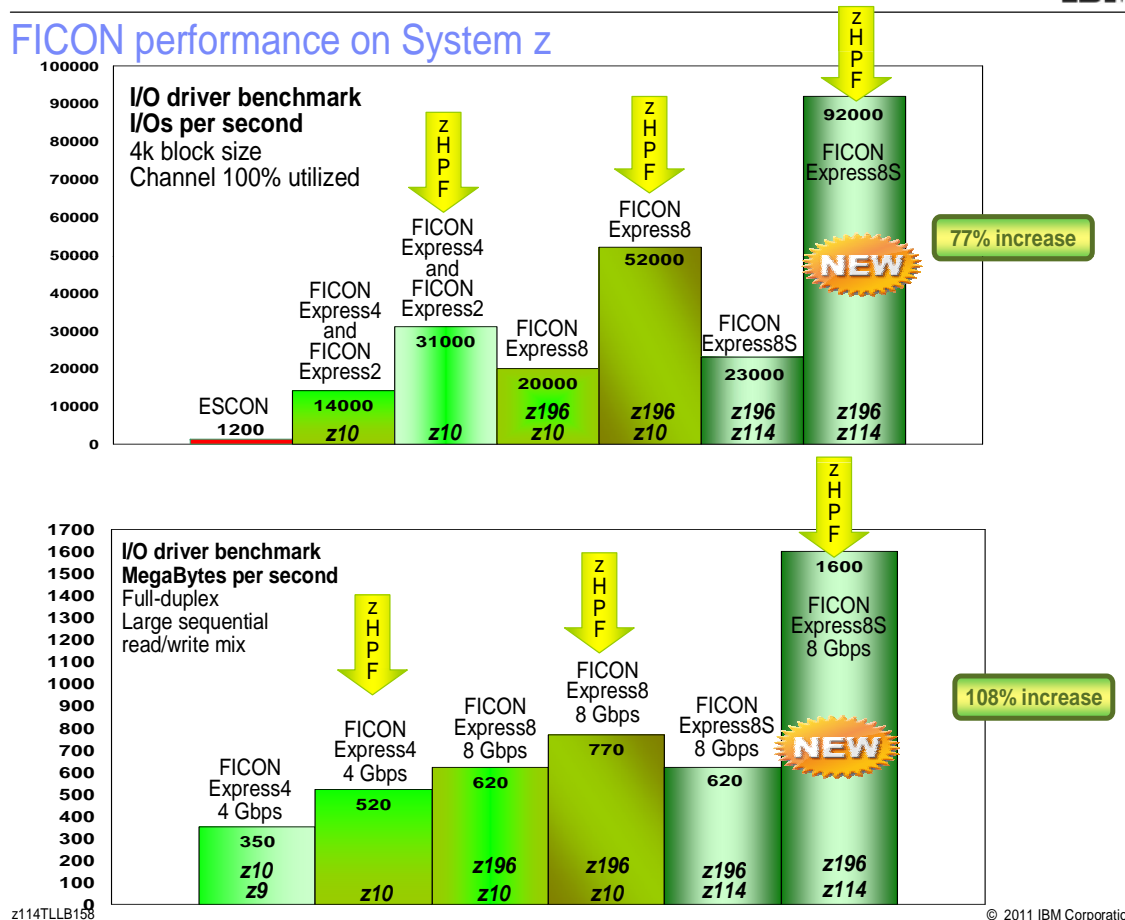
Dr. Steve Guendert
Principal Engineer



Channel consolidation basic considerations

- Initial (late 1990s) rule of thumb (ROT) for aggregating ESCON channels to FICON channels was 8:1
- Since then I/O channel technology has advanced greatly (see next slide).
- Which channels will use zHPF?
- What about the storage connectivity?
- Each scenario today is unique, best practice is to do an analysis.
 - IBM FICON aggregation analysis or similar.
- **Always consider response time and RAS**

Channel performance



Graphic from new System z Channel Migration paper

Some possible ratios

Channel Types	Feature Code (LX)	Feature Code (SX)	Channels per Card	FICON MBps Throughput Performance and zHPF capable		IO/sec Capable with 4K Block Size at 100% Utilized with zHPF	IO/sec Capable with 4K Block Size at 100% Utilized without zHPF	Estimated ESCON to FICON Aggregation Ratios without zHPF
				READS or WRITES Half Duplex	READS + WRITES Full Duplex			
FICON Express8S - 8 Gbps	409	410	2	600 (or 800 w/zHPF)	600 (or 1600 w/zHPF)	92,000	23,000	19:1
FICON Express8 - 8 Gbps	3325	3326	4	600 (or 800 w/zHPF)	600 (or 770 w/zHPF)	52,000	20,000	16:1
FICON Express4 - 4 Gbps	3321/3324/3323	3322/3318	2	330 (or 400 w/zHPF)	350 (or 520 w/zHPF)	31,000	14,000	11:1
FICON Express2 - 2 Gbps	3319	3320	4	200	270	31,000	13,000	10:1
FICON Express - 1 Gbps	2319	2320	2	170	170	N/A	9,200	7:1
ESCON F/C 2324 - 17 MBps	N/A	N/A	15+1 spare	17	N/A	N/A	1,200	N/A

Graphic from new System z Channel Migration paper

Recommended utilization reference lines

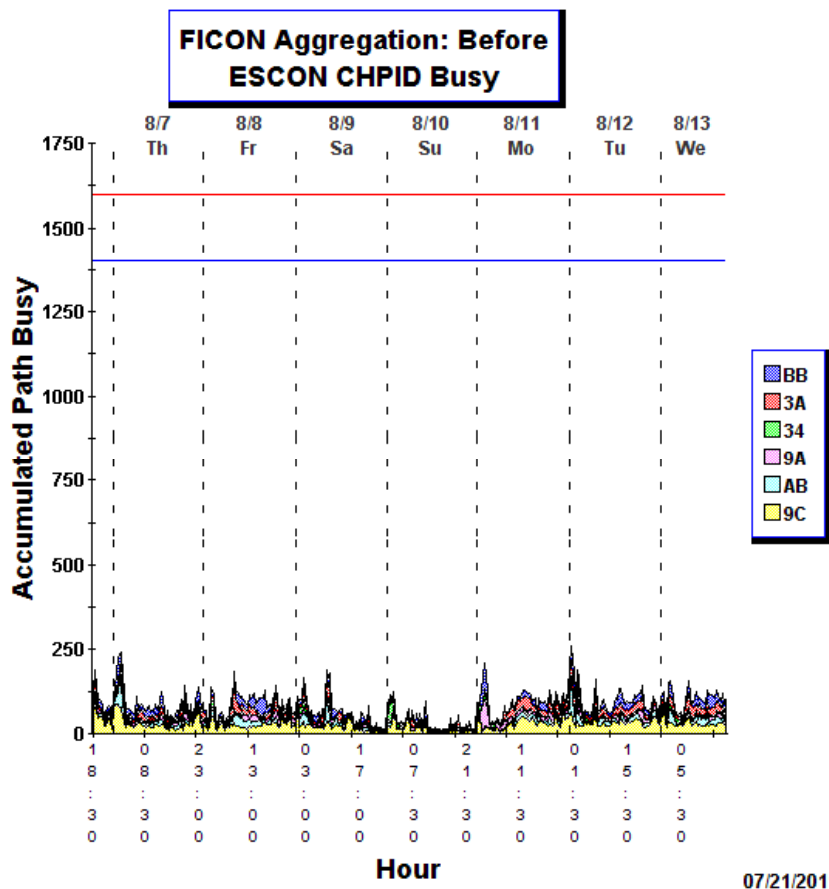
- For FICON Express CHPID utilization, the recommended maximum utilization level is 50%.
- For the FICON Bus busy utilization, the recommended maximum utilization level is 40%.
- For the FICON Express Link utilization with an estimated link throughput of 2Gbps/Sec, 4Gbps/Sec or 8Gbps/Sec, the recommended maximum utilization threshold level is 70%.

SOME TOOLS FOR YOUR TOOLBOX

IBM zCP3000

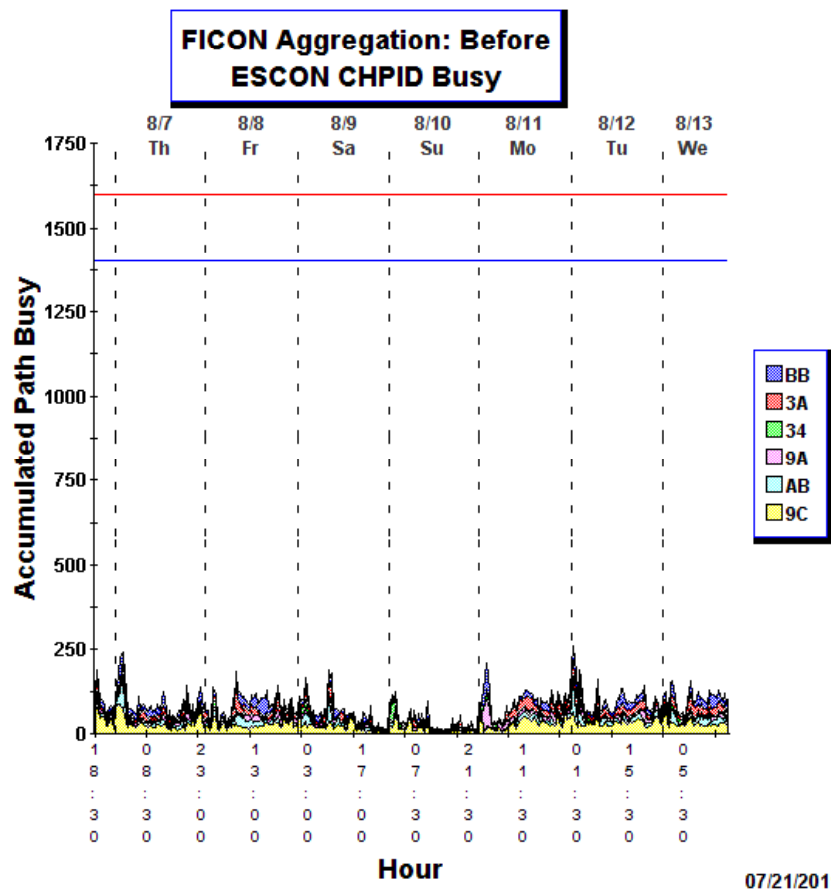
- IBM's capacity planning tool for System z hardware
 - z/OS, z/VM
- Sophisticated tool used by IBM FTSS personnel on behalf of customers-typically at no charge.
 - Incumbent to know your data, and take data that included peak periods
- Examples on slides which follow
- FICON aggregation analysis with zCP3000 should be accompanied with a Disk Magic and/or Tape Magic study to provide a complete channel to device evaluation, which would include response time analysis

zCP3000 example



- This chart above shows before aggregation of 6 ESCONS Channels after aggregation occurs to 1 FICON Express 8GB Channel.
- This area graph shows the sum of measured channel utilization for the ESCON channels selected for a proposed aggregation for the selected paths.
- Each ESCON channel is represented by a separate area within the graph, and the selected ESCON channels are listed to the right of the graph.
- The accumulated utilization for all 6 CHPIDs has a range of [10.1%,263.1%]. The average value is 81.6%. The maximum occurs on 8/12/08 at 00:30.

zCP3000 example cont'd



- The Red line at 1600% for aggregation onto a FICON Express8 channel represents the maximum combined ESCON channel utilization recommended when there will be at least 2 native FICON control unit ports logically daisy-chained to the same FICON Express channel and the workload driving the peak ESCON utilizations is primarily large blocksize batch activity.
- The blue line at 1400% on this chart is more appropriate for those customers whose peak utilizations are driven by short blocksize online database transaction workloads.

zCP3000 example cont'd

CHPID	Type	BCU	Min % Busy	Max % Busy	Avg % Busy	Max occurs at:
BB	E	8400	0.1	76.8	15.6	8/11/08:02:00
3A	E	8000	0.4	84.2	15.0	8/8/08:08:00
34	E	7B05	1.2	65.2	3.5	8/10/08:01:00
9A	E	7A00	3.3	77.8	9.5	8/11/08:01:30
AB	E	7640	1.1	87.2	11.3	8/12/08:00:30
9C	E	7000	3.2	89.1	26.7	8/7/08:01:00
ALL 6	---	---	10.1	263.1	81.6	8/12/08:00:30

- Key shows that after Aggregation of 6 FICON channels, estimate busy on one FICON Express 8GB is less than 5 % percent busy.


zHPF analysis tool

- IBM ATS analysis service (free)
- When properly configured, zHPF is designed to help reduce overhead and improve performance.
- Prior to implementing zHPF customers should, determine the impact of zHPF with their current workload.
- The purpose of the study is to quantify the amount of I/Os in the customer's environment which may be zHPF eligible.
- With this information the impact/benefit of zHPF can be assessed for the customer's specific configuration and workload.


zHPF analysis tool example

- This chart shows with zHPF they can handle about 60% more I/O per sec.
- The Key point Of this chart reveals that the projected channel utilization is less when zHPF is applied

% I/O Eligible for zHPF and % Channel Capacity Increase with zHPF




Current CHPID Utilization Compared to Projected CHPID Utilization with zHPF



Graphic from new System z Channel Migration paper


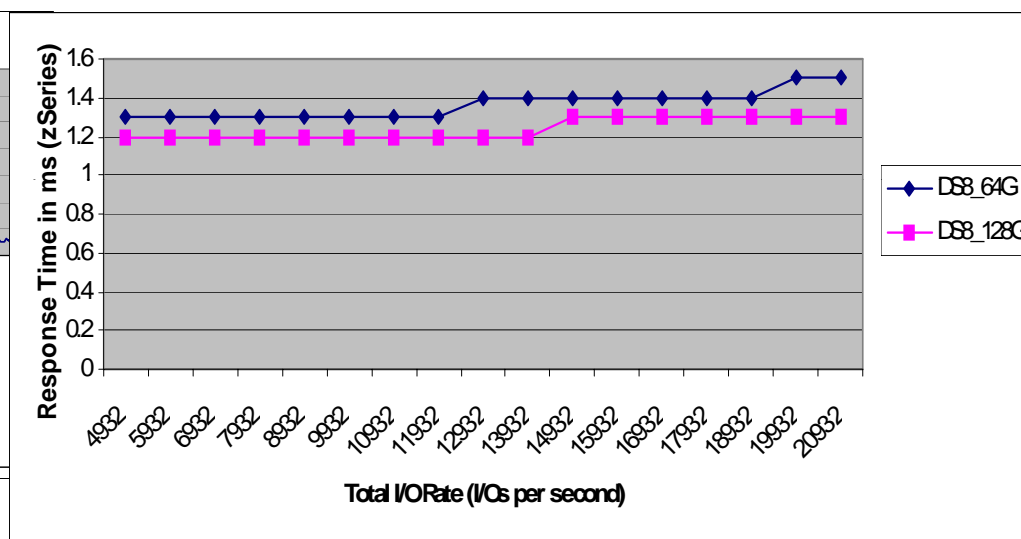
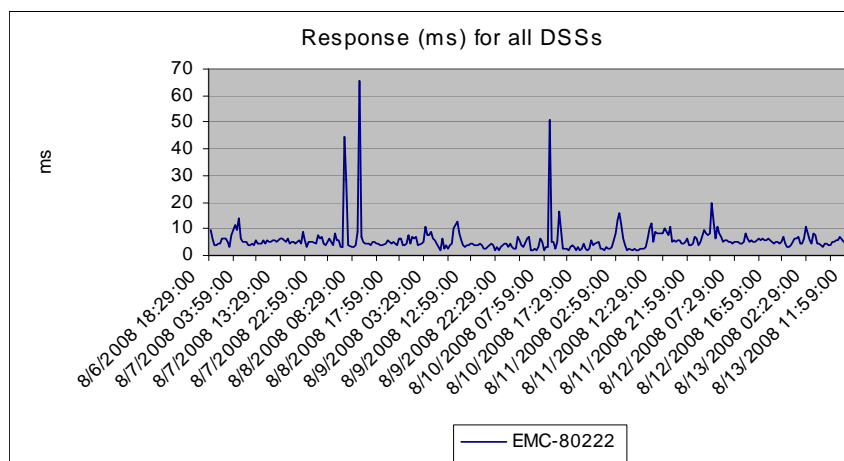
Intellimagic Disk Magic (Direction)

- IBM has licensed Disk Magic from IntelliMagic since 1994.
- Is available to end users as well.
- The end-user version provides support for other disk storage systems including IBM, HP, HDS and EMC. Both IBM and non-IBM storage customers benefit from direct license of the product.
- In July 2010 IntelliMagic changed the name of the end-user version of Disk Magic to IntelliMagic Direction.
- provides a way to model existing workloads and combining new workloads on any major hardware platform.
- End user can do a myriad of “what if” modeling.
 - For example, you can model the impact of migrating from FICON Express 2 to FICON Express 8 and zHPF.
 - You can simultaneously model the impact of implementing Extended Address Volumes, HyperPAVs, etc.



SHARE
Technology · Connections · Results

Disk Magic examples




Intellimagic RMF Magic (Vision)


- In July 2010 RMF Magic was enhanced and renamed to Intellimagic Vision.
- IntelliMagic Vision can be used for the sizing and monitoring of FICON configurations, including those with inter-switch links (ISL).
- With IntelliMagic Vision you can determine the throughput on any channel, ESCON or FICON, and you can easily determine how much bandwidth is needed for each disk subsystem.
- IntelliMagic Vision **provides recommendations on the number of aliases required** for both HyperPAV and Dynamic PAV configurations.
- IntelliMagic Vision also provides bandwidth sizing and analysis capabilities for the asynchronous copy services implementations from all the major System z storage vendors.

RMF Magic examples

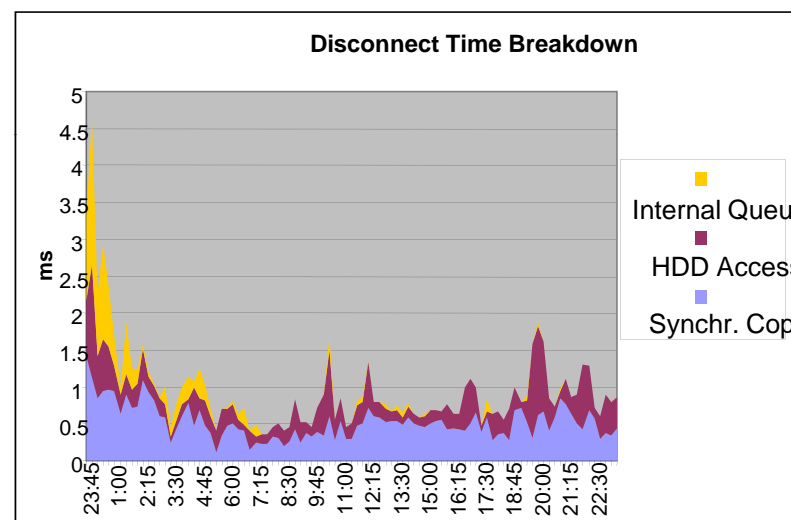
Effective Channel Data Rate (MB/s) for all DSSs



Potential Bandwidth Savings with Asynchronous Copy
Write Bandwidth (MB/s)



Disconnect Time Breakdown



Questions?



Planning for the Elimination of ESCON I/O on System z

August 11, 2011
Session # 9899

