



HITACHI
Inspire the Next

All Flash Array Customer Case Study

Walter Amsler - Senior Director
Global Office of Technology and Planning
Hitachi Data Systems

THE CHALLENGING TRANSITION TO S.M.A.C

SMAC = Social Business, Mobility, Big Data Analytics, Cloud

Source: Capgemini – Key IT Requirements 2014

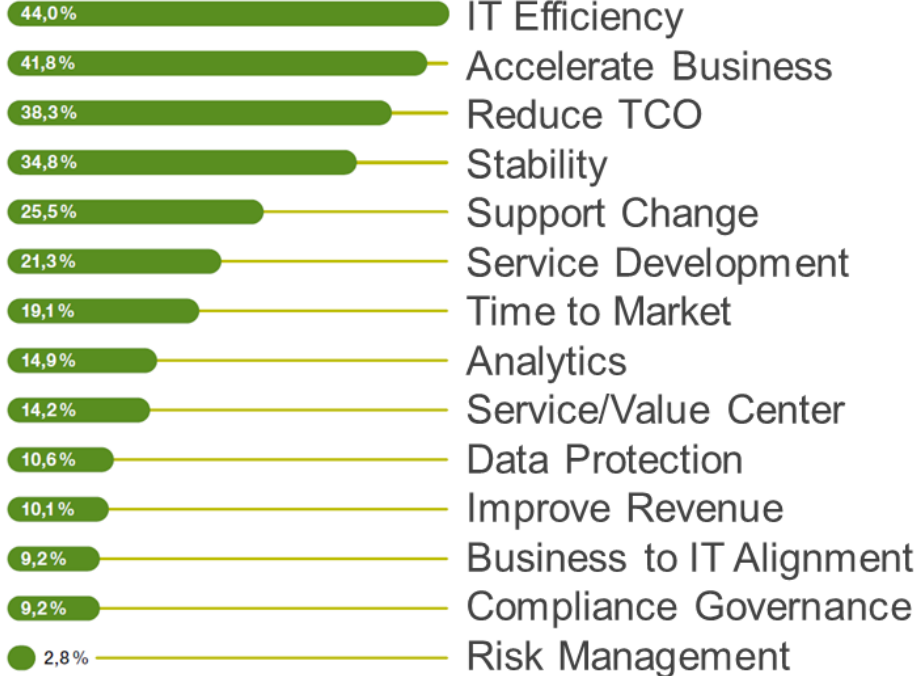


**24x7
Expectations**

**Too much
Complexity**

**Too much
Information**

**Continually
rising Costs**



- The Coop Group is the leading retail company in Switzerland
 - comprises retail companies in Switzerland as well as wholesale and manufacturing companies in Switzerland and abroad
 - \$28,174M CHF Total Sales Revenue in 2014 (\$30B)
 - \$1,124 Net Sales from Online Shops (24.1% YoY Growth)
 - 1971 Retail Sales Outlets in Switzerland, 199 Wholesale and Production Sites
- IT Technology Profile
 - Core business runs on SAP with Oracle, on IBM P-Series Servers with AIX
 - Using SAP/HANA for realtime business intelligence
 - Other Platforms e.g. VMware, Solaris, Linux etc.
 - Long distance (120km) Asynchronous Remote Replication for DR



Issues and Requirements



■ Challenges

- Current Storage solution unable to keep up with increased demands
 - nightly Business Warehouse consolidation of distribution center data
 - online retail application YoY Growth
- Insufficient visibility into storage health - difficult to diagnose performance issues and size latent demand
- limited headcount – rising complexity – increased business requirements

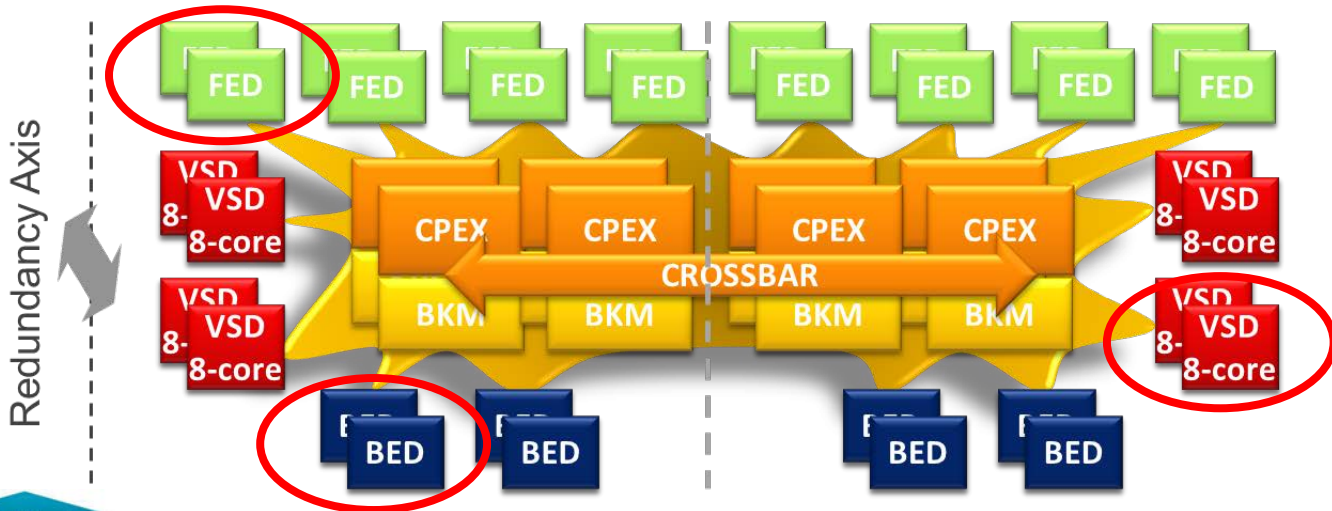
■ Requirements

- Always On Infrastructure
- Investment Protection and predictable cost for Hardware and Software purchases - future proof technology
- All Flash strategy to mitigate performance bottlenecks and reduce complexity

The solution: Hitachi VSP G1000

All Flash Array with Universal Replicator


128 x 8/16Gbps SAN Ports




Rear View



Front View

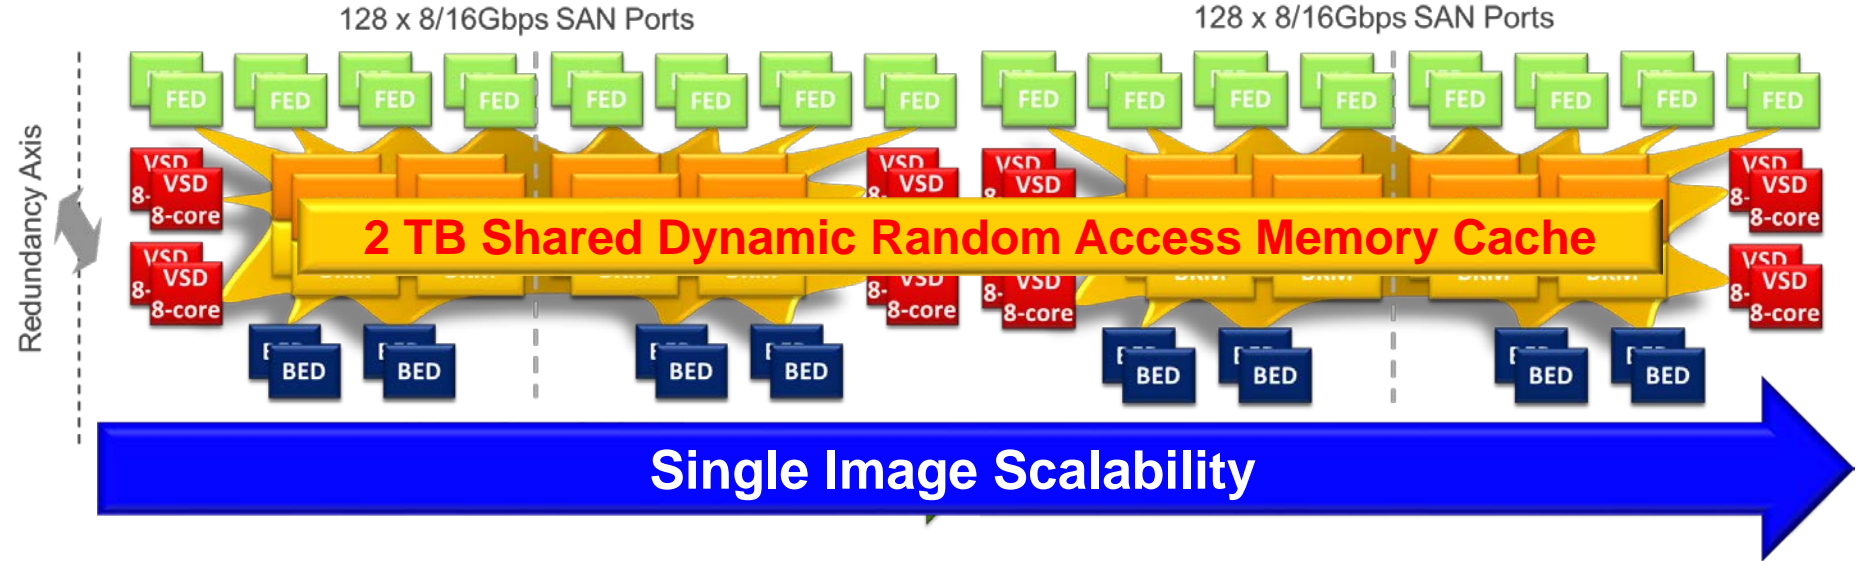


TB Capacity – 134 * 3.2TB Hitachi Flash Modules



Purpose built and optimized for Flash

VSP G1000 Hyper scale-up architecture



Future Proof Technology - predictable cost defined for 1.6PB Capacity

Future Flash technologies will be easily integrated into the Storage virtualization OS e.g. 3D NAND, TLC, denser Flash Modules, offload engines for data reduction etc.

Remote Replication for DR – 2 Data Centers


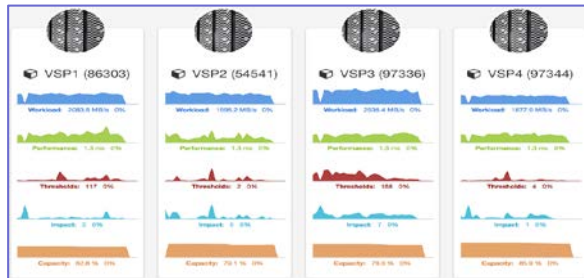
- Asynchronous – no distance limitations
 - CPU Offload - reduced cost by replacing host based replication
 - Supports virtually any platform, operating system and application
 - No Host Impact on Response Times
 - Data Movement Overhead on Out-Of Region Array (Pull Technology)
 - Journaling maintains update sequence Integrity
 - RPO measured in seconds, RTO in minutes (application dependent)






Hitachi Data Systems IOportal Cloud Service - EoU

- Dashboard
 - KPIs
 - Management View
 - Mobile Devices
- Capacity
 - Aggregation on Site, Storage System and Pool
 - Tier mixing Ratio
 - Detailed Cluster (Server) Reports
- Performance
 - Server Workload
 - Storage System Performance
 - Worktime








Production IO workload



HUR 120 km



HDS VSP G1000 S/N 58270 (vsp1bs) Top10 Port timeseries for entire storage system.
The entire storage system has 44 connected ports.
Average interval 300 sec, UTC offset 0 sec. Graph created by user Henry_Schmidt on Mon::27-Jul-2015_14:14:06



I/O Response Time

SVCT[ms] — Low@2 ms — Medium@10 ms — High@100 ms



Top 10 Servers



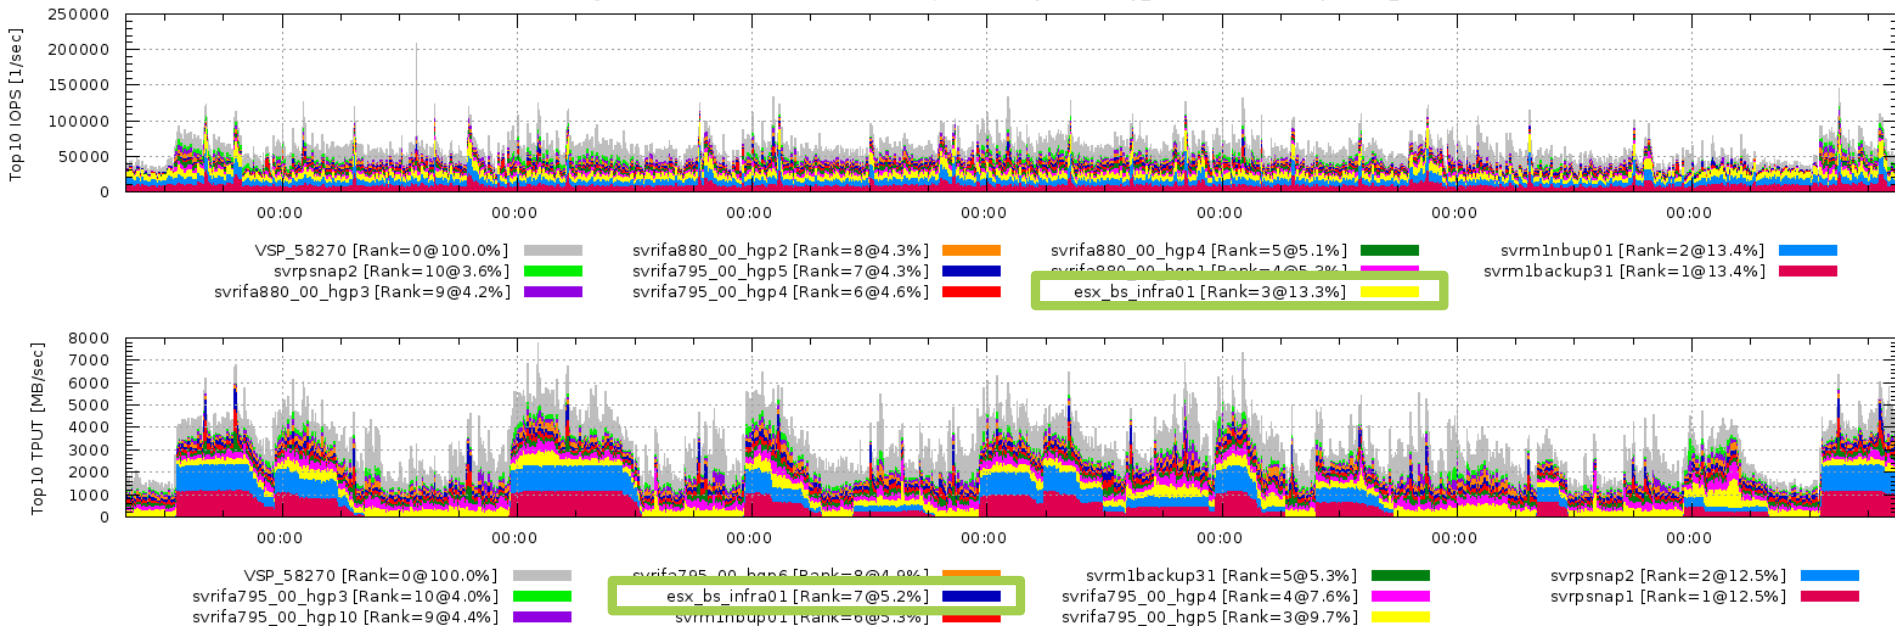
ESX Server is ranked 3 for IOPS, ranked 7 for Throughput

HDS VSP G1000 S/N 58270 (vsp1bs) Top10 Server timeseries performance for entire storage system.

The entire storage system has 3624 LUNs and 36 active Cluster-Interfaces: CL1-B CL1-C CL1-D CL1-E CL1-F CL1-K CL1-L CL1-M CL1-N CL1-P CL2-B CL2-C CL2-D CL2-E CL2-F CL2-K CL2-L CL2-M CL2-N CL2-P CL3-A CL3-B CL3-C CL3-D CL3-J CL3-K CL3-L CL3-M CL4-A CL4-B CL4-C CL4-D CL4-J CL4-K CL4-L CL4-M






The entire storage system has 57 connected servers (i.e. Host Storage Domains).

Average interval 300 sec, UTC offset 0 sec. Graph created by user Henry_Schmidt on Wed::29-Jul-2015_09:24:44



ESX Server KPI's

HDS VSP G1000 S/N 58270 (vsp1bs) LU Server timeseries performance for esx_bs_infra01. Rank 3
Server esx_bs_infra01 has 130 LUNs and 4 active Cluster-Interfaces: CL1-C CL1-L CL2-C CL2-L
Average interval 300 sec, UTC offset 0 sec. Graph created by user Henry_Schmidt on Mon::27-Jul-2015_14:14:19



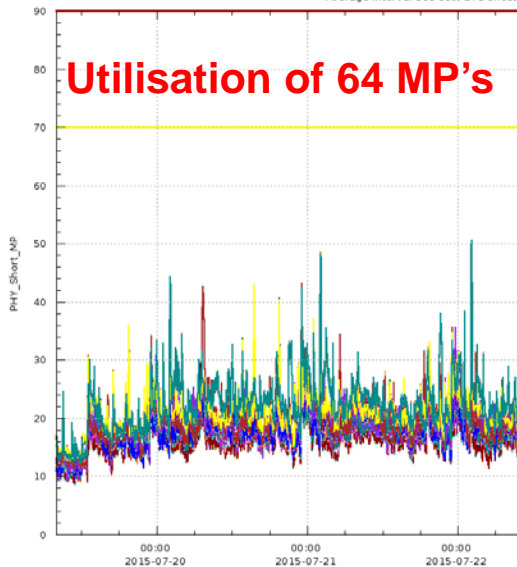


Many More Metrics are available e.g. MP Busy and Remote Replication

HDS VSP G1000 S/N 58270
Res

Font: Deja
Average interval 300 sec. UTC offset

Utilisation of 64 MP's



- pqVspStdPerf.awk V4.9.2 by peaq
- MPB-2MD.MP00-2MD [Rank=64@1.32%]
 - MPB-2MD.MP01-2MD [Rank=63@1.32%]
 - MPB-2MD.MP03-2MD [Rank=62@1.33%]
 - MPB-2MD.MP05-2MD [Rank=61@1.33%]
 - MPB-2MD.MP04-2MD [Rank=60@1.33%]
 - MPB-2MD.MP06-2MD [Rank=59@1.33%]
 - MPB-2MD.MP02-2MD [Rank=58@1.33%]
 - MPB-2MD.MP07-2MD [Rank=57@1.33%]
 - MPB-1MA.MP00-1MA [Rank=56@1.42%]
 - MPB-2MC.MP01-2MC [Rank=55@1.43%]
 - MPB-2MC.MP00-2MC [Rank=54@1.43%]
 - MPB-1MA.MP01-1MA [Rank=53@1.43%]
 - MPB-2MC.MP05-2MC [Rank=52@1.43%]
 - MPB-2MC.MP03-2MC [Rank=51@1.43%]
 - MPB-2MC.MP02-2MC [Rank=50@1.44%]
 - MPB-2MC.MP06-2MC [Rank=49@1.44%]
 - MPB-1MA.MP05-1MA [Rank=48@1.44%]
 - MPB-2MC.MP04-2MC [Rank=4
 - MPB-1MA.MP03-1MA [Rank=4
 - MPB-2MC.MP07-2MC [Rank=4
 - MPB-1MA.MP06-1MA [Rank=4
 - MPB-1MA.MP07-1MA [Rank=4
 - MPB-1MA.MP02-1MA [Rank=4
 - MPB-1MD.MP04-1MD [Rank=4
 - MPB-1MD.MP00-1MD [Rank=4
 - MPB-1MD.MP01-1MD [Rank=2
 - MPB-1MD.MP05-1MD [Rank=3
 - MPB-1MD.MP03-1MD [Rank=3
 - MPB-1MD.MP06-1MD [Rank=3
 - MPB-1MD.MP07-1MD [Rank=3
 - MPB-1MD.MP02-1MD [Rank=3
 - MPB-2MA.MP00-2MA [Rank=3
 - MPB-2MA.MP01-2MA [Rank=3


HDS VSP G1000 S/N 58270 (vsp1bs) UniversalReplicator timeseries performance for metric UniversalReplicator.

Resource UniversalReplicator, metric UniversalReplicator has 12 instances.



Synopsis: No synopsis for L1=UniversalReplicator, L2=UniversalReplicator found in database

Font: DejaVuSans-Bold.ttf, font size: 10, color theme: standard.

Average interval 300 sec. UTC offset 0 sec. Graph created by user Henry_Schmidt on Wed:29-Jul-2015_09:24:38




> 1GB/sec Write data Replication Throughput




Longterm Trending Data (April 1st – July 29th)

- Proactive performance management leads into capacity management and problem avoidance


IOPS




Troughput



Response Time




Blocksize



- All Flash Array with Enterprise Class Functionality and High Availability
 - Array based asynchronous replication for any Server, OS and application, removes the need for Host Software and server cycles and therefore reduces SW and HW cost
 - Average Response Times well below 1 millisecond (ms) for entire storage subsystem
 - Daily sequential throughput peak > 6GB/second
- The need for Millions of IOPS – Myth or Reality?
 - Even in a highly consolidated environment with a variety of OS (AIX, VMware, Linux, Solaris) and a diversity of applications (BI, ERP, Analytics), average I/O Rate observed is around 50,000 I/O's per second, with peaks of 120,000 IOPS
 - Blocksize is massively increasing, average Blksize is 40kB and up to 80kB for reads
 - Most important is the very consistent low response times – High IOPS rating indicates subsystem scalability without impact to low latency performance
- Need storage subsystem performance metrics for deep insights
 - The right Metrics and instrumentation help deliver Quality of Service –Visibility!

- New subsystem coped with substantial latent demand
 - Eliminated hours of Elapsed Time for critical applications
 - Reduced Response Time by factor of 10
- Impact on Web Shop application
 - Faster shopping cart checkout time
 - Improved Online customer satisfaction
- Impact on Coop Business Warehouse
 - daily consolidation of distribution centre data now completes in time
 - Can do a lot more analytics runs, allows better planning and decision making, with SAP/HANA can do realtime stock adjustment planning



End User Quote: Habt Ihr über das Wochenende SAP "frisirt"? Geht alles viel schneller...konnte nicht mal mehr ins Kaffee...Die Auswertungen waren schon auf dem Bildschirm bevor ich aufstehen konnte...


VSP G1000 AFA performance leadership Storage Performance Council (SPC)

SPC1
#1



OLTP Workload
Highly Random
Highly Parallel
High IOPS / Low Latency

**2,004,941 IOPS @
0.96ms Response Time**



Business Relevant

SPC2
#2



DSS Workload
Highly Sequential
Highly Parallel
High Throughput

43,012 MB/s