

.NEXT CONFERENCE

— Business Critical Apps on Nutanix / 고객사례 : 신세계 아이앤씨

Josh Odgers, Solutions Architect/

강승근 팀장, 신세계 아이앤씨

.NEXT
ON TOUR

.NEXT CONFERENCE

— Scalability, Resiliency & Performance.

The perfect platform for business critical applications

Josh Odgers
Staff Solution Architect
NPX#001 / VCDX#090

.NEXT
ON TOUR

— Introductions

- Josh Odgers @josh_odgers

Staff Solutions Architect

Double-VCDX (DCV / Cloud)

Nutanix Platform Expert (NPX#001)



.NEXT
ON TOUR

— Agenda

The three pillars of success for Infrastructure / Business Critical Apps



Why traditional infrastructure is NOT the answer



How Nutanix delivers the three pillars of success



Business Critical Application deployment example - Shinsegae



Q & A



POP QUIZ:

What percentage of Nutanix deployments are for Enterprise Applications?



BONUS QUESTION: What percentage of nodes run AHV?

— Agenda

The three pillars of success for Infrastructure / Business Critical Apps



Why traditional infrastructure is NOT the answer



How Nutanix delivers the three pillars of success



Business Critical Application deployment example - Shinsegae



Q & A



— 3 Pillars of success for Infrastructure / Business Critical Apps



— Agenda

Resiliency

- Data Integrity
- Failure scenarios
 - SSD / HDD
 - Node / Server
 - Controller
 - Multiple failures
- Self Healing
- Management
- Physical workloads

Scalability

- CPU / RAM
- Capacity
- Performance
- Management
- Scale up & out
- Physical workloads

Performance

- IOPS
- Throughput
- Noisy Neighbor
- Mixed Workloads
- During Failures
- Management
- Scale up & scale out apps
- Physical workloads

— Agenda

The three pillars of success for Infrastructure / Business Critical Apps



Why traditional infrastructure is NOT the answer



How Nutanix delivers the three pillars of success



Business Critical Application deployment example - Shinsegae




Q & A




Storage Performance Scaling on SAN

Day 1 – Performance is good! 😊




Storage Performance Scaling on SAN

I.T Manager “We need to support more users, let’s add another server”




Storage Performance Scaling on SAN

I.T Manager: “Performance isn’t getting any better! Let’s add some more servers.”



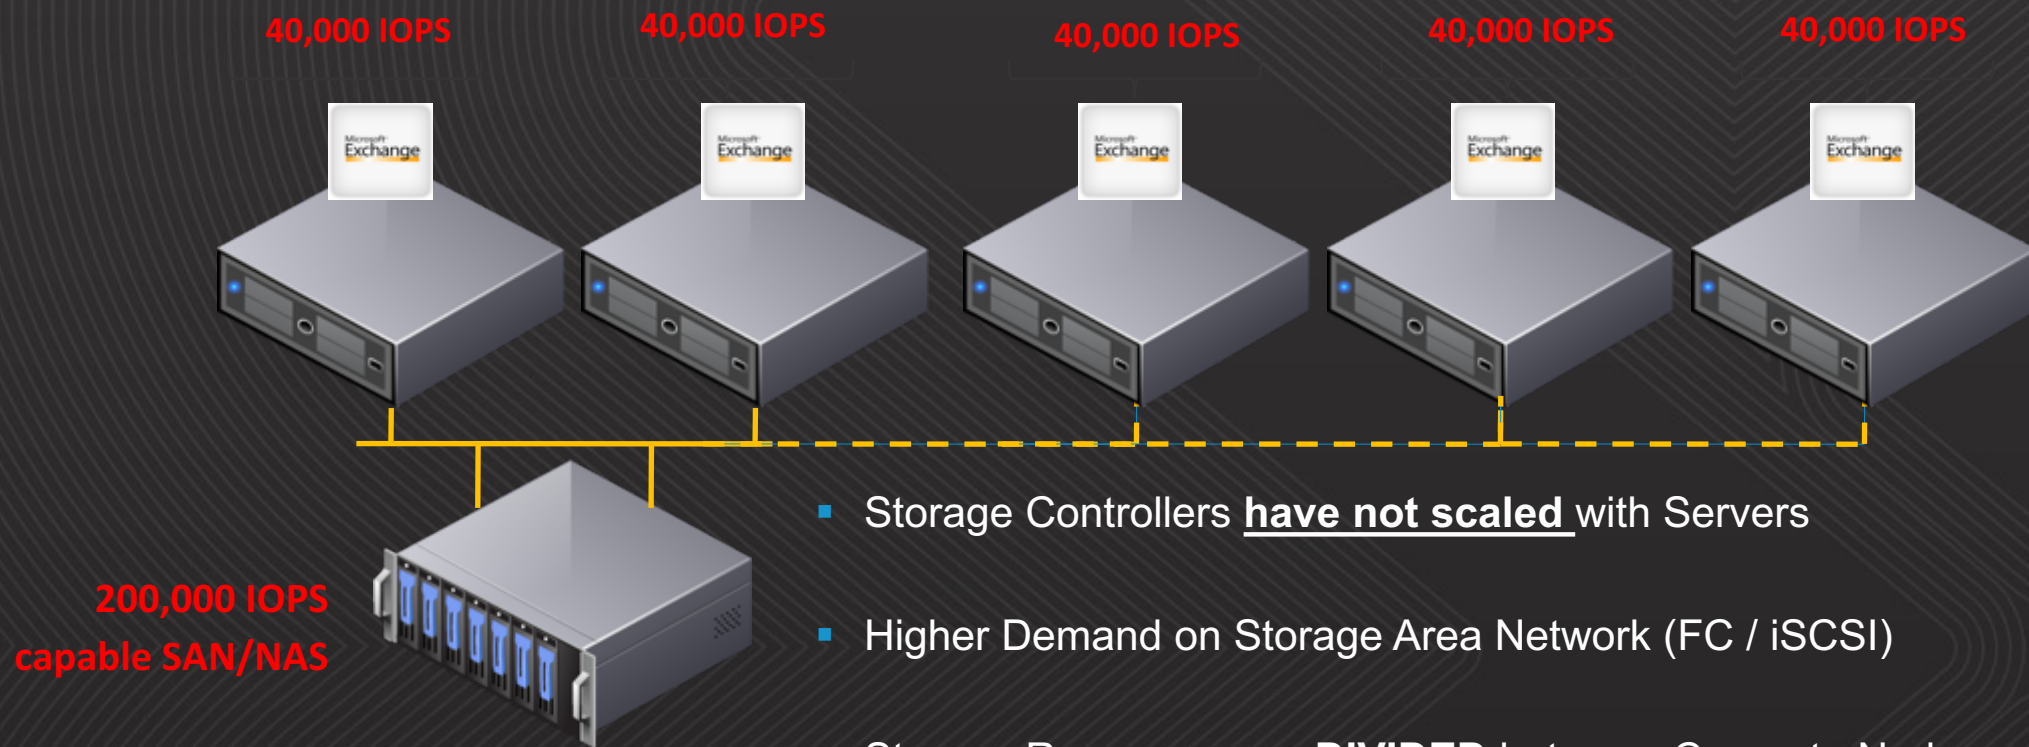
Storage Performance Scaling on SAN

I.T Manager: “Why is our performance getting worse, we just invested more money into servers”



Storage Performance Scaling on SAN


I.T Admin “The performance problems continue getting worse. Why?”



- Storage Controllers have not scaled with Servers
- Higher Demand on Storage Area Network (FC / iSCSI)
- Storage Resources are DIVIDED between Compute Nodes
- Compute nodes COMPETE for storage performance with other nodes even when on different LUNs

Storage Performance Scaling on SAN

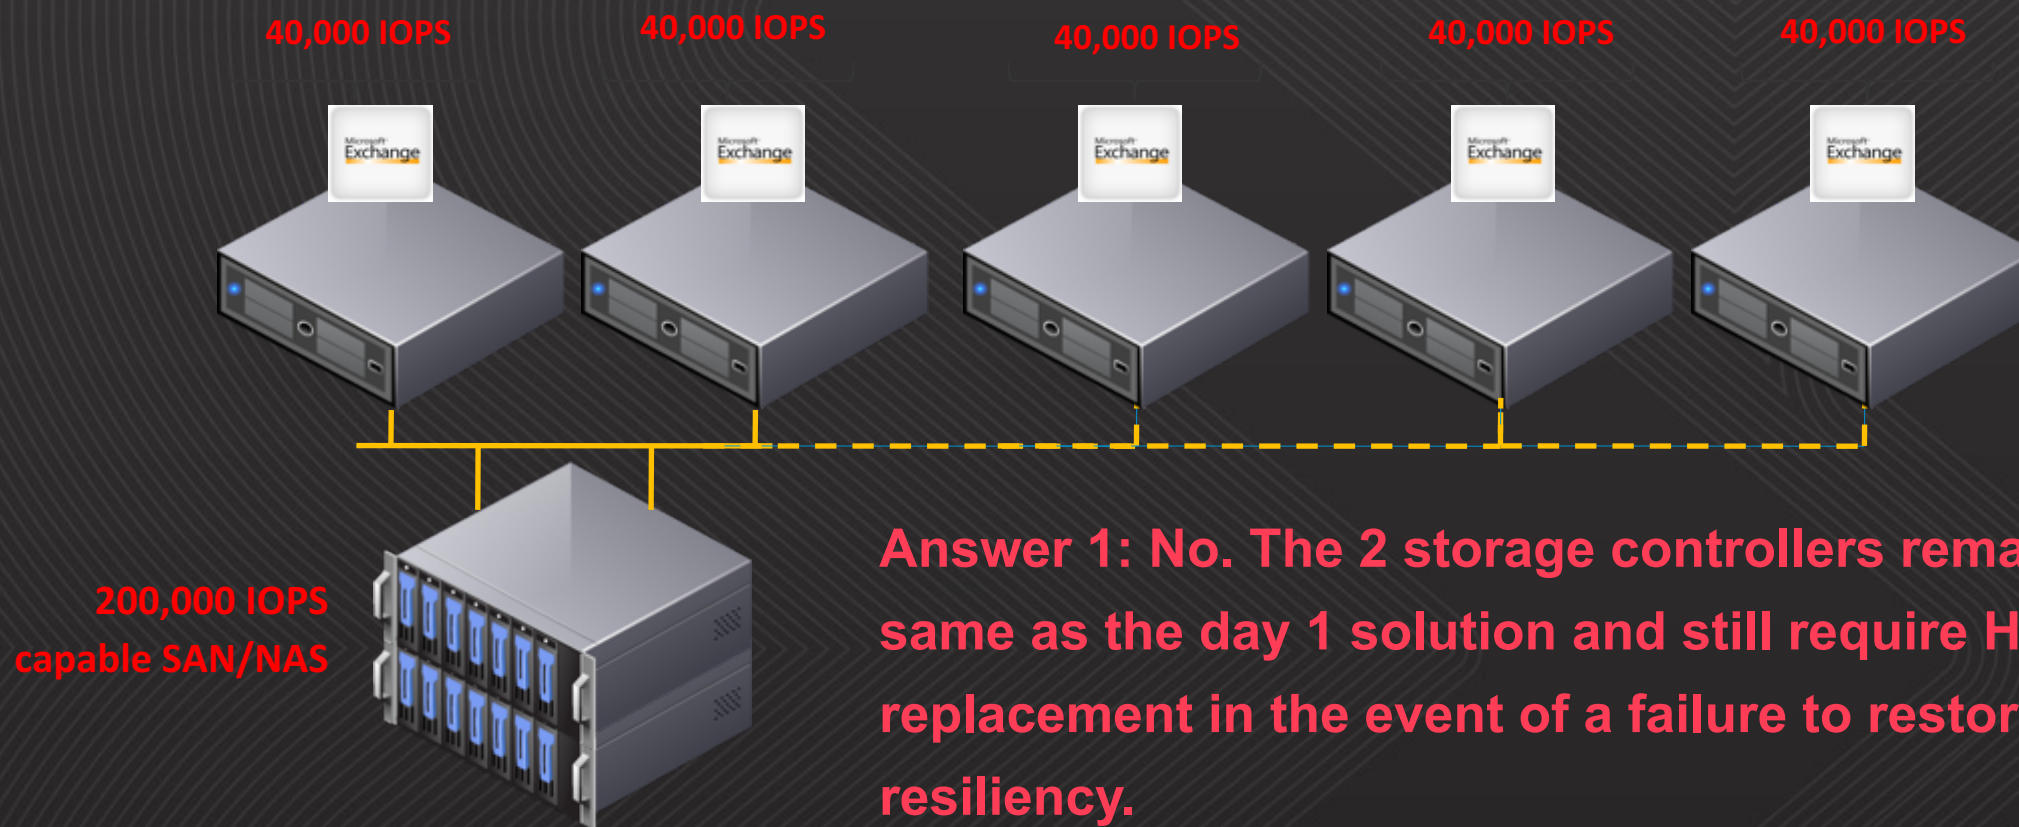
Performance has reduced again! A.K.A CV generating event!



What about adding disk shelves?

QUESTION ONE

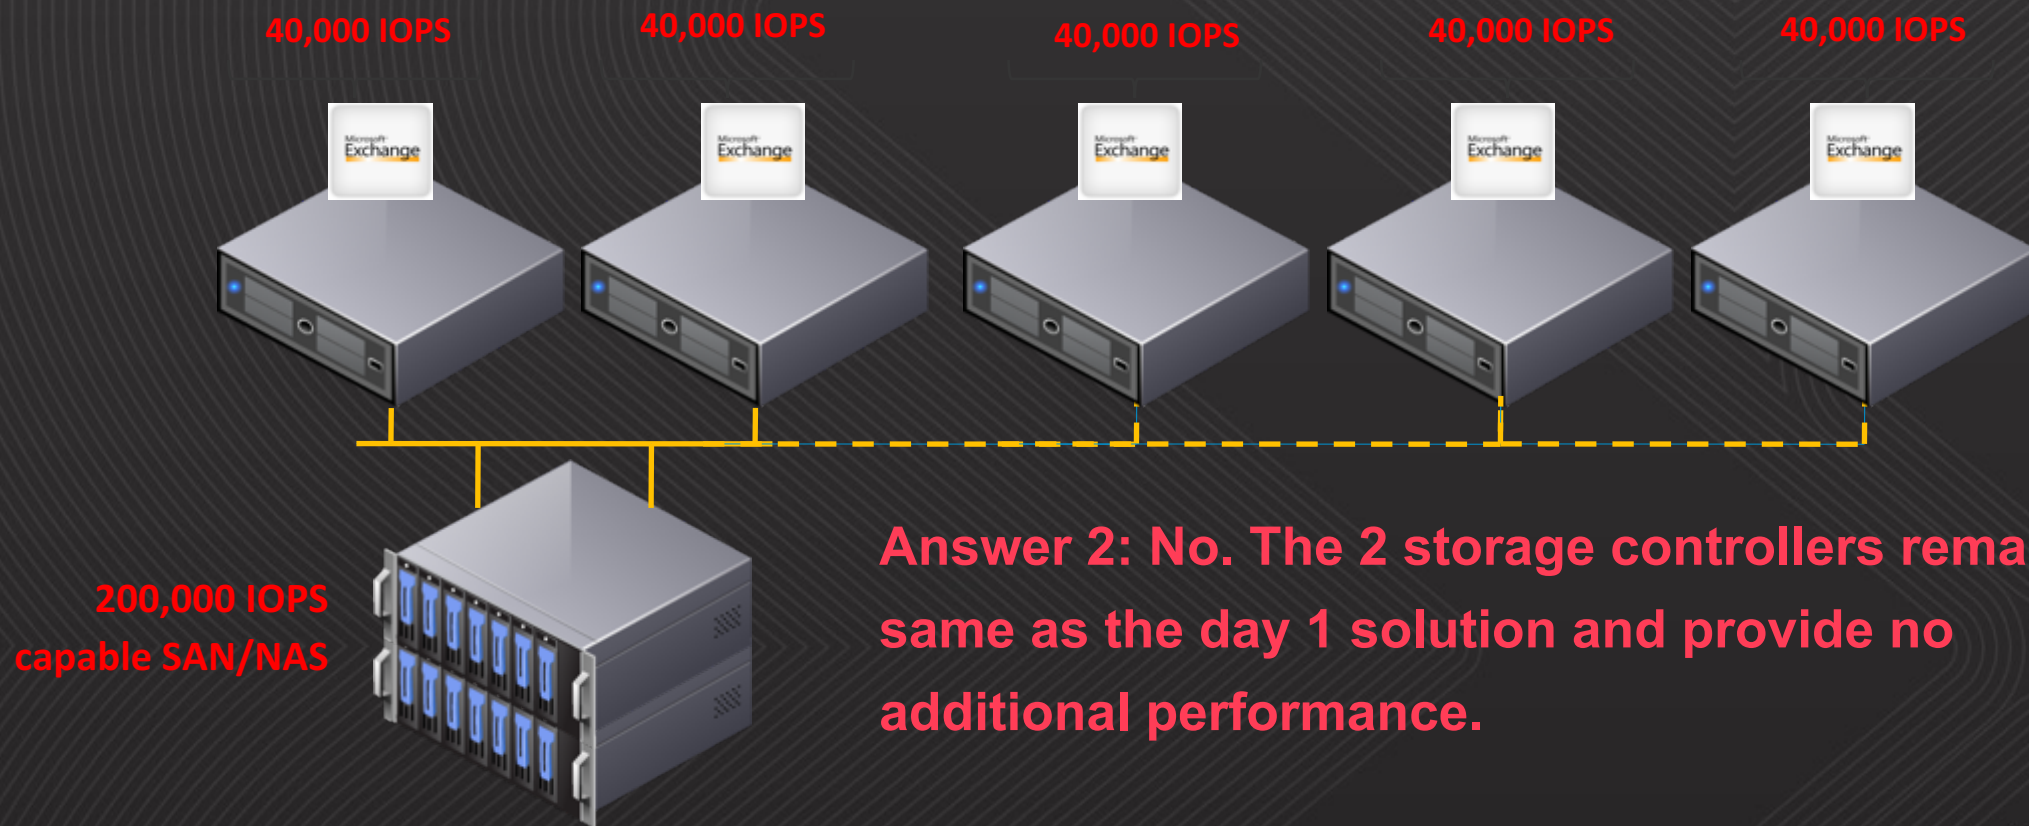
Question 1: Have we improved storage resiliency since Day 1?



Answer 1: No. The 2 storage controllers remain the same as the day 1 solution and still require HW replacement in the event of a failure to restore resiliency.

QUESTION TWO

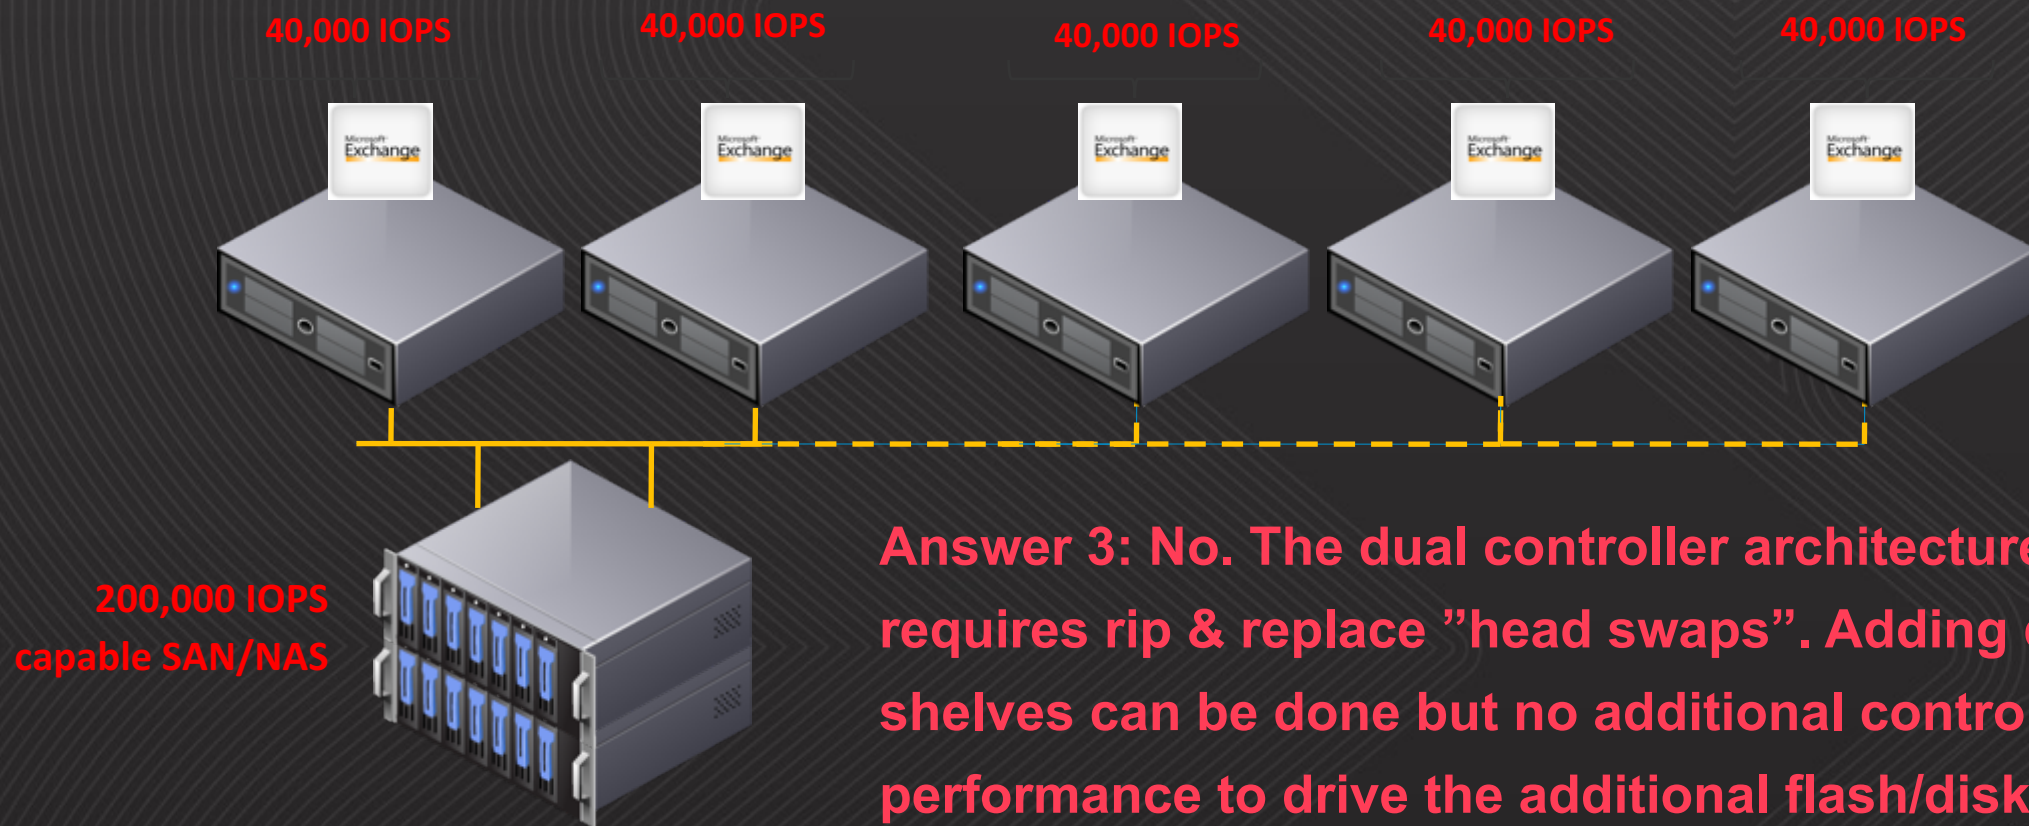
Question 2: Have we improved storage performance since day 1?



Answer 2: No. The 2 storage controllers remain the same as the day 1 solution and provide no additional performance.

QUESTION THREE

Question 3: Is a traditional 3-tier solution easily scalable?



Answer 3: No. The dual controller architecture requires rip & replace "head swaps". Adding disk shelves can be done but no additional controller performance to drive the additional flash/disk.

QUESTION:

Does/Can traditional 3-tier infrastructure deliver the three pillars of success?



— Agenda

The three pillars of success for Infrastructure / Business Critical Apps



Why traditional infrastructure is NOT the answer



How Nutanix delivers the three pillars of success



Business Critical Application deployment example - Shinsegae




Q & A




— How Nutanix delivers the three pillars of success - [Resiliency](#)

The underlying Nutanix distributed architecture since Day 1, is the key!




- 5 Node Cluster functioning normally

— How Nutanix delivers the three pillars of success - [Resiliency](#)




- 5 Node Cluster functioning normally
- One node fails – vSphere HA restarts VM on a surviving host
- VM functions normally – accesses data locally/remotely as required

— How Nutanix delivers the three pillars of success - [Resiliency](#)




- 5 Node Cluster functioning normally
- One node fails – vSphere HA restarts VM on a surviving host
- VM functions normally – accesses data locally/remotely as required
- Nutanix detects some data does not have two copies & repairs cluster

— How Nutanix delivers the three pillars of success - [Resiliency](#)




- 5 Node Cluster functioning normally
- One node fails – vSphere HA restarts VM on a surviving host
- VM functions normally – accesses data locally/remotely as required
- Nutanix detects some data does not have two copies & repairs cluster
- VM data is localized over time as data is read over the 10G LAN

— How Nutanix delivers the three pillars of success - [Resiliency](#)



- 5 Node Cluster functioning normally
- One node fails – vSphere HA restarts VM on a surviving host
- VM functions normally – accesses data locally/remotely as required
- Nutanix detects some data does not have two copies & repairs cluster
- VM data is localized over time as data is read over the 10G LAN
- Cluster is fully functional and fully repaired
- Only 20% of the controllers are lost (vs 50% for legacy storage)

— How Nutanix delivers the three pillars of success - [Resiliency](#)




- A 2nd host fails - vSphere HA restarts VM on a surviving host
- VM functions normally – accesses data locally/remotely as required
- Nutanix detects some data does not have two copies & repairs cluster
- VM data is localized over time as data is read over the 10G LAN
- Cluster is fully functional and fully repaired
- AND can tolerate a further controller failure!
- Only 40% of the controllers are lost (vs 100% for legacy storage)

— How Nutanix delivers the three pillars of success - [Resiliency](#)

Key points about Nutanix self healing / resiliency / recoverability.

- Not Constrained by traditional RAID
- Rebuild performance not limited to a single hot spare
- A single hot spare slows down restoration of resiliency
- The slower the recovery the....
 - Higher the impact on performance
 - Higher risk of subsequent failure
 - Higher risk of downtime and/or data loss
- **Nutanix ADSF rebuilds:**
 - Are driven by all controller VMs (CVM) in a cluster
 - Read and write from all drives, to all drives.
 - Low impact thanks to distributed storage fabric
 - Minimize the chance of subsequent failure
 - DO NOT require SSD/HDD to be replaced to restore resiliency & tolerate further failures


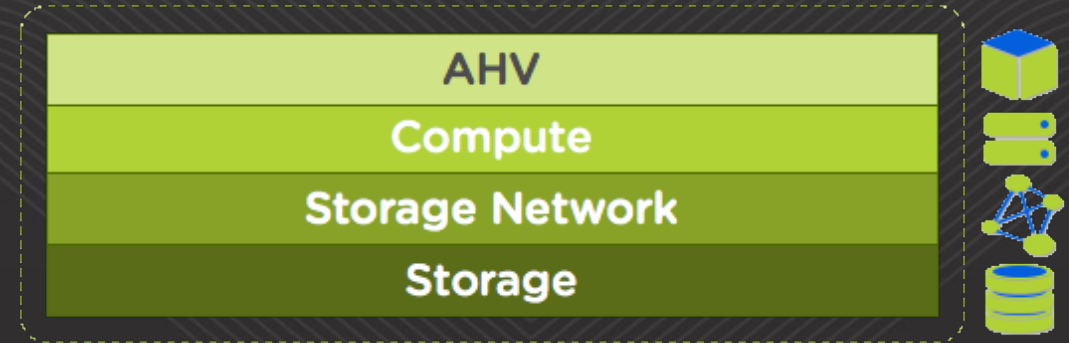


— How Nutanix delivers the three pillars of success - [Resiliency](#)

Nutanix Resiliency, includes our distributed management: PRISM

▪ Key points about Nutanix management plane resiliency

- PRISM Element is built into every cluster
- PRISM Element scales automatically with the cluster
- The Management layer resiliency is automatically recovered after failures



— How Nutanix delivers the three pillars of success - [Resiliency](#)

Summary:

Drive failures are transparent to applications

Fully automated / Low Impact self healing

Proactive monitoring / analytics / planning

Ability to self heal without HW replacement

Management stack is also highly available



Example 1: With RF2 and block awareness, Nutanix can support the loss of up to 24 drives concurrently!




Example 2: With RF3 and block awareness, Nutanix can support the loss of up to 48 drives concurrently!



— How Nutanix delivers the three pillars of success - [Resiliency](#)

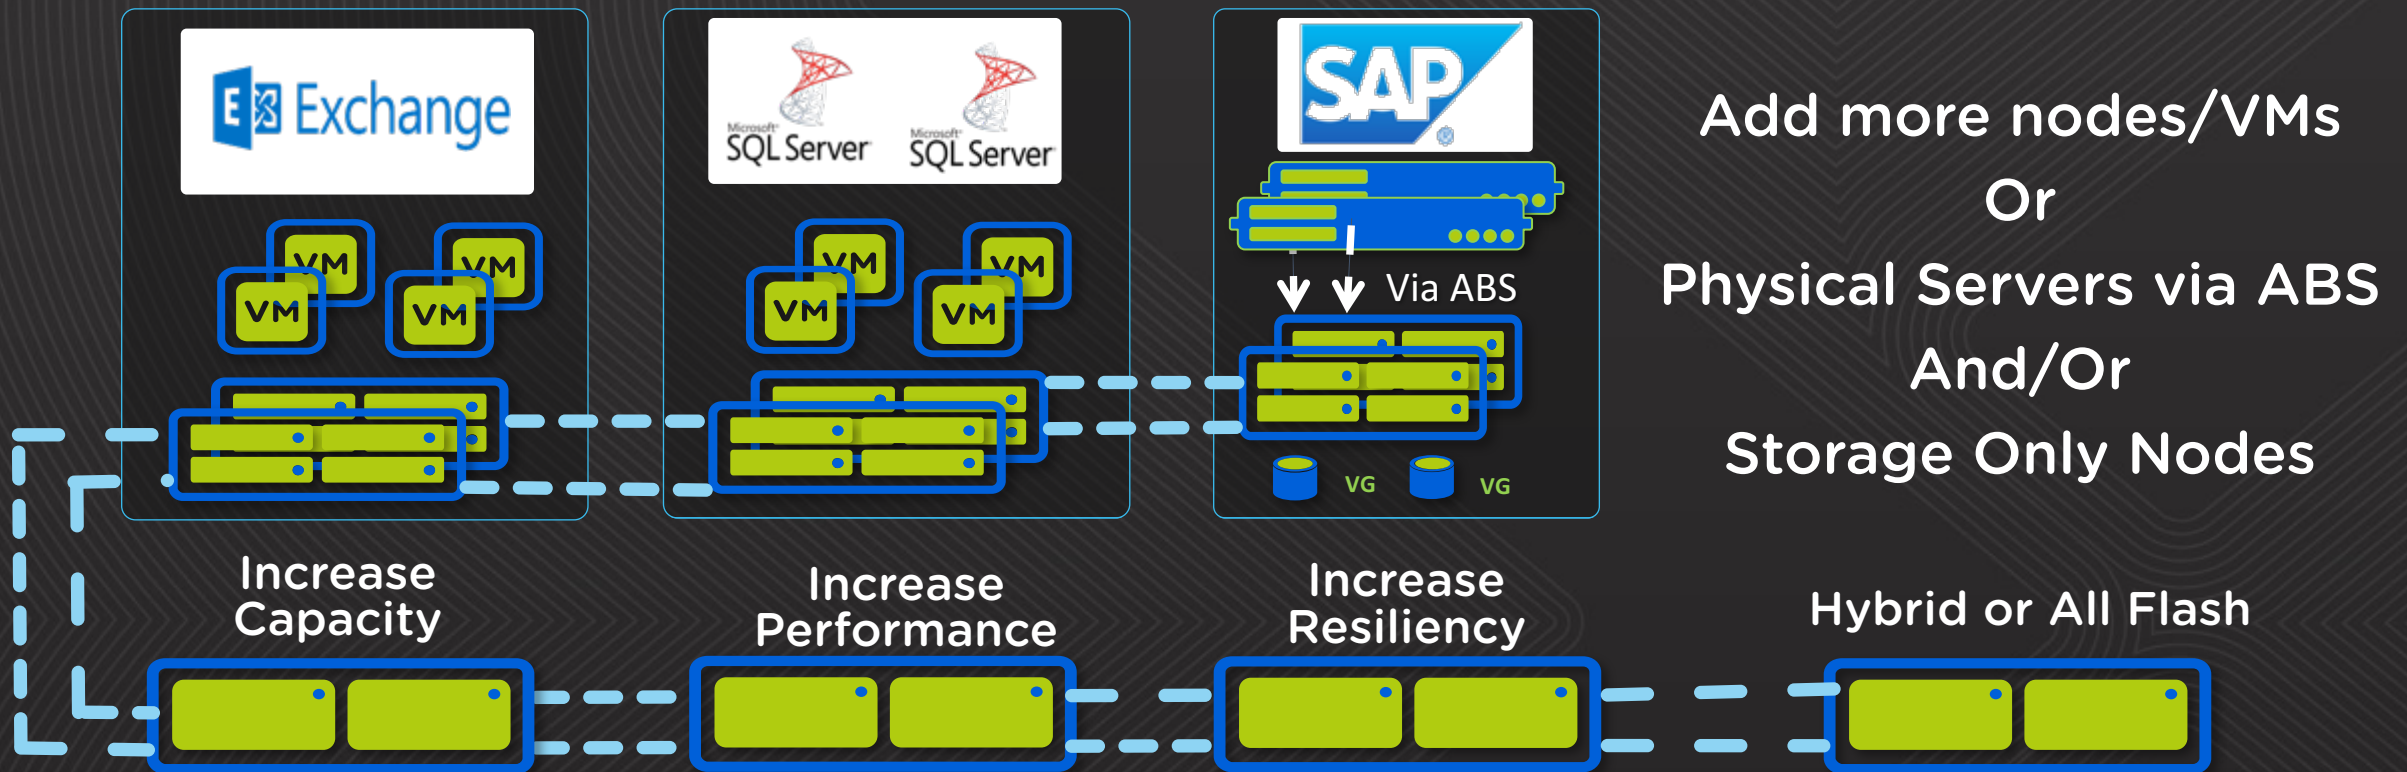
Nutanix Resiliency protects.....



— How Nutanix delivers the three pillars of success - [Scalability](#)

Example Nutanix environment scaling to support many use cases

- Mix node types & Hardware generations
- Physical servers connected via ABS can utilize the entire clusters capacity/performance
- Mix All Flash and Hybrid nodes
- Scale linearly or non-linearly as required




Nutanix scales while adding capacity, resiliency & performance!

— How Nutanix delivers the three pillars of success - [Scalability](#)


Example: Physical server connected to a 4 node Nutanix cluster via Acropolis Block Services (ABS)

Nutanix cluster expanded from 4 to 8 nodes, the number of Active Paths for IO increases automatically.

4 Node Cluster = 4 Active Paths


Network Activity  3 Gbps Network I/O		
Image	PID	Address ^
System	4	192.168.1.31
System	4	192.168.1.32
System	4	192.168.1.33
System	4	192.168.1.34

8 Node Cluster = 8 Active Paths

Network Activity  1052 Mbps Network I/O		
Image	PID	Address ^
System	4	192.168.1.255
System	4	192.168.1.31
System	4	192.168.1.32
System	4	192.168.1.33
System	4	192.168.1.34
System	4	192.168.1.61
System	4	192.168.1.62
System	4	192.168.1.63
System	4	192.168.1.64


— How Nutanix delivers the three pillars of success - [Scalability](#)

Nutanix Scalability delivers increased:



— How Nutanix delivers the three pillars of success - [Scalability](#)





Example: 32 Node cluster, 4 different Models, multiple different HW generations

 **Josh Odgers** @josh_odgers · Jun 11

Scalability/Flexibility are key to I.T infrastructure.
#Nutanix provides excellent scalability & supports mixed HW,node types etc.
#HCI #CIO

Hardware Summary


32 HOSTS	14 BLOCKS	NX-3450	1 Block
		NX-3050	2 Blocks
		NX-6020	4 Blocks
		NX-3060-G4	1 Block
		NX-6035-G4	4 Blocks
		NX-8035-G4	2 Blocks

  12  14 

—How Nutanix delivers the three pillars of success - [Performance](#)

Key requirements for real world performance.


- Ability to scale
 - Non disruptively
 - Without impacting resiliency
 - While INCREASING
 - Resiliency
 - Performance
 - Capacity



—How Nutanix delivers the three pillars of success - [Performance](#)

Data Locality


- Local reads avoid having network bandwidth, congestion becoming a bottleneck
- Flash is now equal to and soon to be exceeding the speeds capable of even 100GB networking
- Reading I/O locally delivers the most CONSISTENT performance
- Writing one replica locally, and one remotely effectively removes 50% of the (write) I/O from the network



—How Nutanix delivers the three pillars of success - [Performance](#)

Scale out shared nothing architecture

- Every node provides additional performance to the cluster thanks to the Controller VM (CVM)
- ADSF intelligent replica placement ensures hotspots are not created in the cluster
- All nodes/drives/CVMs in the cluster contribute to read and write activity



—How Nutanix delivers the three pillars of success - [Performance](#)

Scale out with Compute+Storage OR storage only nodes

- Where more VM (CPU/RAM) performance is required, standard Compute+Storage nodes add performance
- For more storage IOPS/Throughput, Storage only nodes can be used

Example: 4 Node Cluster with 4 Jetstress VMs with 4 Nodes

	Jetstress01	Jetstress02	Jetstress03	Jetstress04	Average across VMs
Performance					
Transactional IOPS Achieved	1063.169	1116.578	1145.858	1165.643	1122.812
Average Latency					
Log Writes	1.478	1.53325	1.578	1.54825	1.53
Database Writes	5.6045	5.88125	5.80475	5.814	5.77
Database Reads	10.047	9.71225	9.569	9.41975	9.68
Initialization					
Database 1 Creation Time	2:23:10	2:26:00	2:26:23	2:24:10	2:24:56
Duplicating 3 Databases Time	4:28:31	4:16:06	4:11:07	4:13:31	4:17:19
Checksum Validation					
DB checksums duration	0:38:31	0:37:16	0:37:48	0:38:28	0:38:01

—How Nutanix delivers the three pillars of success - [Performance](#)

Scale out with Compute+Storage OR storage only nodes

- Where more VM (CPU/RAM) performance is required, standard Compute+Storage nodes add performance
- For more storage IOPS/Throughput, Storage only nodes can be used

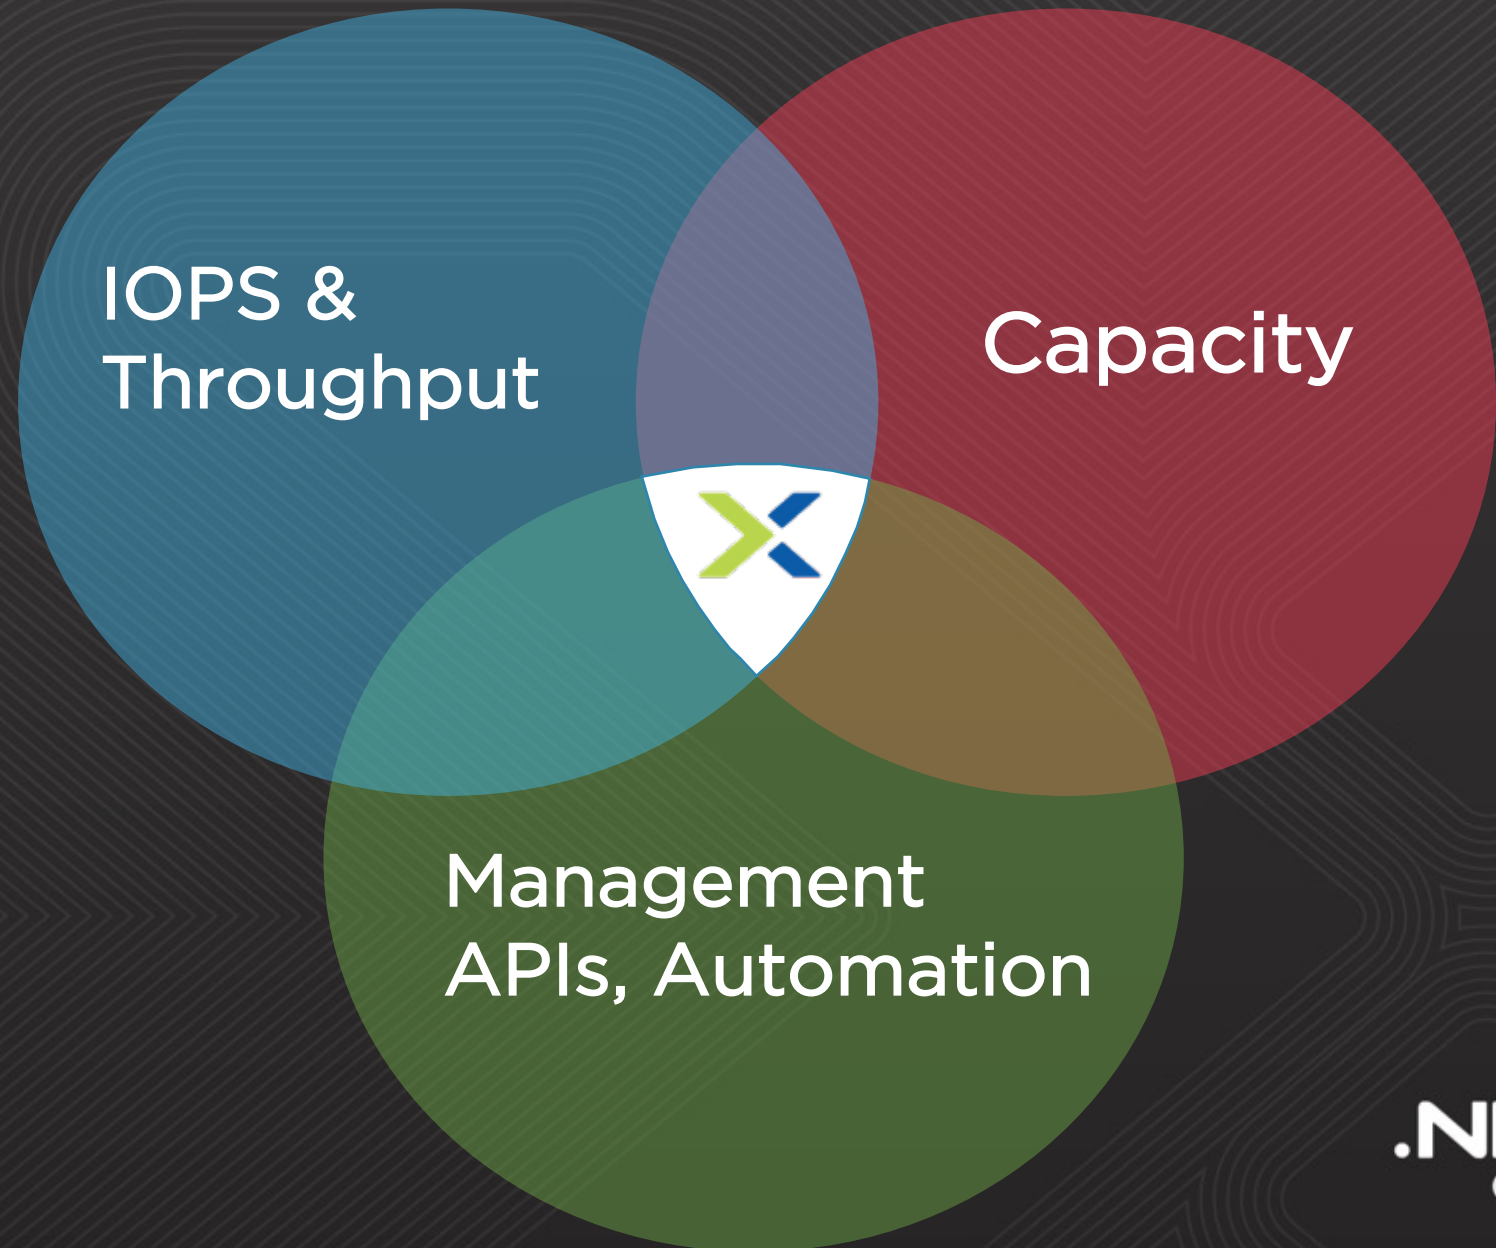
Example: 4 Node Cluster with 4 Jetstress VMs with 4 Nodes + [4 Storage only nodes](#)

	Jetstress01	Jetstress02	Jetstress03	Jetstress04	Average across VMs
Performance					
Transactional IOPS Achieved	2167.805	2142.476	2216.038	2217.978	2186.07 (+194.70%)
Average Latency					
Log Writes	1.34975	1.373	1.3515	1.353	1.35 (-13.09%)
Database Writes	4.8	4.79975	4.70675	4.745	4.76 (-21.27%)
Database Reads	5.51925	5.43025	5.169	5.251	5.34 (-81.32%)
Initialization					
Database 1 Creation Time	2:07:07	2:12:01	2:11:36	2:09:17	2:10:00 (-0:14:55)
Duplicating 3 Databases Time	3:04:48	2:57:46	2:57:49	2:58:54	2:59:49 (-0:34:54)
Checksum Validation					
DB checksums duration	0:37:21	0:37:24	0:37:23	0:37:13	0:37:20 (-0:00:40)

[2x the IOPS, Lower Latency across the board, faster DB creation times.](#)


— How Nutanix delivers the three pillars of success - [Performance](#)

Nutanix delivers scalable
performance for:



QUESTION:

Does Nutanix
infrastructure
deliver the three
pillars of success?



— Agenda

The three pillars of success for Infrastructure / Business Critical Apps



Why traditional infrastructure is NOT the answer



How Nutanix delivers the three pillars of success



Business Critical Application deployment example - Shinsegae



Q & A



.NEXT CONFERENCE

— 신세계그룹 협업시스템 Nutanix 도입사례

강승근 팀장 / 신세계아이앤씨

01. 신세계그룹 현황

Retail

Department Stores
Hypermarkets
Supermarkets
Convenience Stores
Mobile Shopping
TV Shopping
Shopping Malls
Premium Outlets
Duty Free

Food

Restaurants
Coffee Chain
Bakeries
Food Manufacturing,
Liquor and Beverage
Distribution

Construction & IT Services

Real Estate, Engineering &
Construction,
IT Services


Hospitality & Leisure

Hotels
GolfClub

Fashion&LifeStyle

Brands
Select Shops

20 Companys + 50,000 Users



— 02. 신세계그룹 협업시스템 “BLOSSOM” & Nutanix



— 02. 신세계그룹 협업시스템 “BLOSSOM” & Nutanix

업무·소통·협업을
하나의 공간에서

Unified
Workplace

민첩,
정확한
커뮤니케이션

Mobile
Workplace

사용자 경험,
업무 흐름을
고려한
디자인(UI/UX)

User-oriented
Workplace

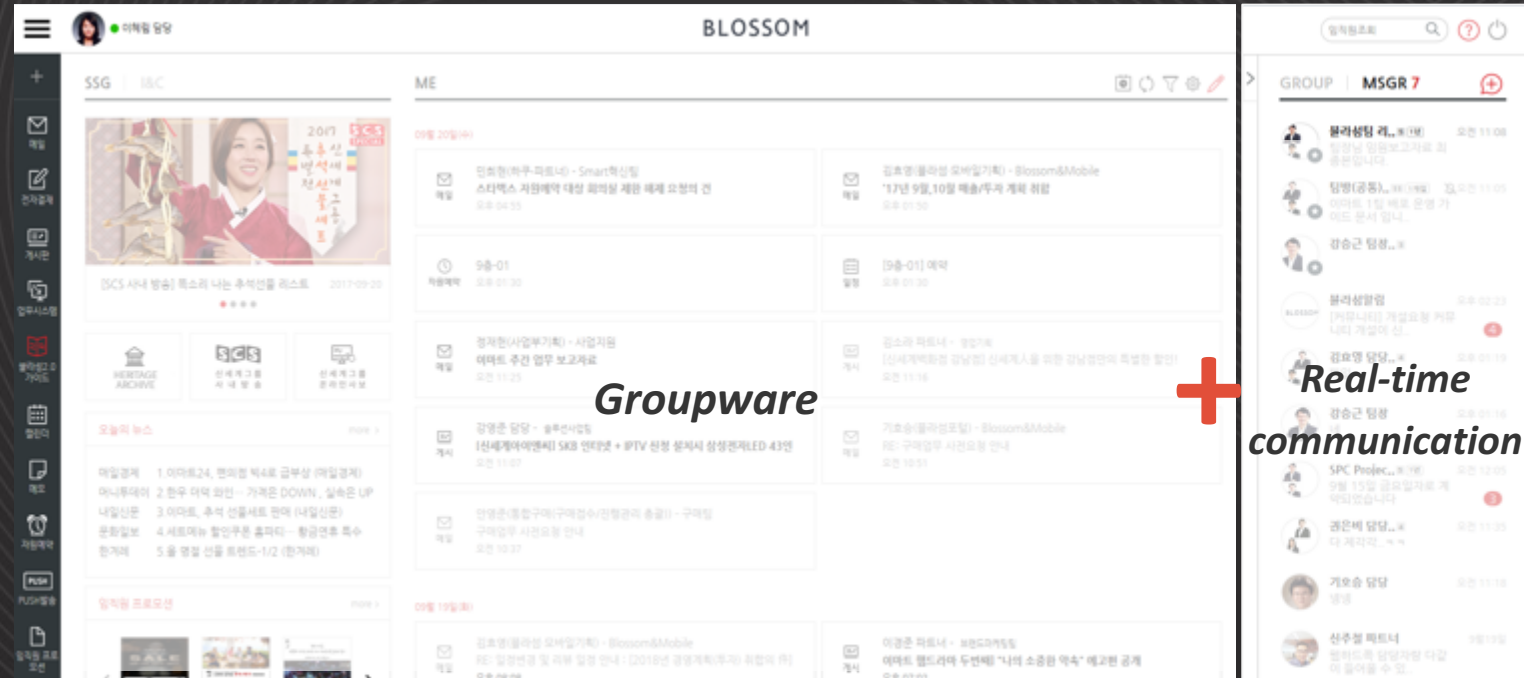
신세계그룹 Digital Workplace 『BLOSSOM 2.0』

02. 신세계그룹 협업시스템 “BLOSSOM” & Nutanix

The image displays the BLOSSOM system interface on two devices. The tablet shows the main dashboard with sections for 'SSG | I&C', 'ME', and 'GROUP | MSGR 7'. The smartphone shows the 'MSGR' (messaging) interface. A large red plus sign is placed between the two devices, with the text 'Groupware + Real-time communication' overlaid.

Groupware + Real-time communication

02. 신세계그룹 협업시스템 “BLOSSOM” & Nutanix



Groupware



Real-time
communication

IIS, MS Exchange, MS-SQL

MongoDB, node.js,
Redis, Angular js

뉴타닉스 솔루션

03. IT인프라 현황 및 개선

As-Is

다양한 시스템 환경

- 멀티 밴더, 멀티 플랫폼 환경
- 펌웨어 업그레이드 어려움
- 성능저하 및 장애시 원인 파악 시간 소요

중앙 집중식 스토리지 환경

- 대부분 시스템이 siloed 환경
- 서버 증설시 스토리지 성능 부하 발생
- 각 분야별 전문가 필요
- 스토리지 용량 증설 시 어려움

비용 증가

- 복잡도에 따른 운영 Cost 증가
- 가상화 라이선스 추가 필요
- 하드웨어 유지보수비 증가

To-Be

다양한 업무시스템을 수용 가능한 단일 플랫폼

- 5만명 MS Exchange Mail
- 다양한 SQL, Oracle Exa 등

운영환경 단순화

- 각 분야별 운영 전문가 최소화
- 성능/용량 계획 용이함
- 확장 유연성

가동율 ↑

총 소요비용 감소 ↓

뉴타니스 솔루션

단일 표준 플랫폼

- MS Exchange / MS SQL
- Oracle RAC
- Windows & Linux workloads

HCI(Hyper-Converged Infrastructure)

- Scale out 구조 제공
- 데이터 분산 저장

One-Click 업그레이드

- BIOS/FW/AOS/Hypervisor 등
순쉬운 업그레이드

손쉬운 유지보수 및 비용절감


- 단일밴더, 단일 플랫폼으로 유지보수 용이
- 오픈소스 기반의 무상 가상화
솔루션(AHV)

04. Nutanix로의 전환

Step 1 3+3 (6) nodes per site (주센터 및 DR)

- Initial Deployment supported 50k User MS Exchange solution

Performance ?



MS Exchange 구성



SCSIC-IB2MBX10 - Pass

- Transactional I/O Performance – 8 Volume 제공 가능 IOPS :1045

Transactional I/O Performance

MSExchange Database ==> Instances	I/O Database Reads Average Latency (msec)	I/O Database Writes Average Latency (msec)	I/O Database Reads/sec	I/O Database Writes/sec	I/O Database Reads Average Bytes	I/O Database Writes Average Bytes	I/O Log Reads Average Latency (msec)	I/O Log Writes Average Latency (msec)	I/O Log Reads/sec	I/O Log Writes/sec	I/O Log Reads Average Bytes	I/O Log Writes Average Bytes
Instance4696.1	5.377	11.082	82.645	47.824	33165.960	35748.892	0.498	2.446	0.090	11.432	3798.031	12794.498
Instance4696.2	5.145	10.891	82.809	48.028	33160.367	35738.716	0.486	2.441	0.091	11.470	3795.180	12769.025
Instance4696.3	5.131	11.149	82.774	47.878	33175.739	35763.910	0.505	2.390	0.091	11.408	3798.031	12733.069
Instance4696.4	5.153	11.066	82.673	47.890	33169.155	35752.917	0.485	2.382	0.091	11.447	3819.417	12746.888
Instance4696.5	5.114	12.038	82.715	47.988	33171.832	35748.152	0.492	2.405	0.091	11.478	3802.308	12775.150
Instance4696.6	5.121	11.775	82.634	47.847	33166.197	35759.704	0.479	2.404	0.090	11.434	3772.369	12869.319
Instance4696.7	5.127	12.040	82.888	47.997	33175.973	35738.771	0.491	2.438	0.091	11.439	3778.072	12682.611
Instance4696.8	5.138	11.838	82.721	47.991	33170.470	35762.897	0.486	2.438	0.090	11.467	3783.774	12812.485

Background Database Maintenance I/O Performance

MSExchange Database ==> Instances	Database Maintenance IO Reads/sec	Database Maintenance IO Reads Average Bytes
Instance4696.1	8.950	261275.503
Instance4696.2	8.963	261227.221
Instance4696.3	8.952	261208.770
Instance4696.4	8.950	261203.266
Instance4696.5	8.953	261175.753
Instance4696.6	8.953	261253.800
Instance4696.7	8.953	261241.336
Instance4696.8	8.952	261235.799

Exchange 2013 Storage Calculator에서 요구하는 수치의 2배를 상회함

04. Nutanix로의 전환

Step 1 3+3 (6) nodes per site (주센터 및 DR)

- Initial Deployment supported 50k User MS Exchange solution

Step 2 6+6 (12) nodes per site (주센터 및 DR)

- Subsequently expanded to 12 at Production hosting MS SQL and Web Servers



Single Nutanix Fabric Across All Workloads

강승근 팀장 / 신세계아이앤씨

— Agenda

The three pillars of success for Infrastructure / Business Critical Apps



Why traditional infrastructure is NOT the answer



How Nutanix delivers the three pillars of success



Business Critical Application deployment example - Shinsegae



Q & A



