



A *Global Payments* Company

WE ARE TSYS

NSX Cluster Solution

User Experience

Roy L. Nicholas

NonStop Technical Boot Camp – 2023



Agenda

1. **TSYS Issuing overview**
2. **NSX cluster solution**
3. **Migration phases 1, 2 and 3**
4. **Dollars and sense**
5. **Summary – Q/A**



PAYMENT STACK CAPABILITIES

Technology
that keeps
you ahead

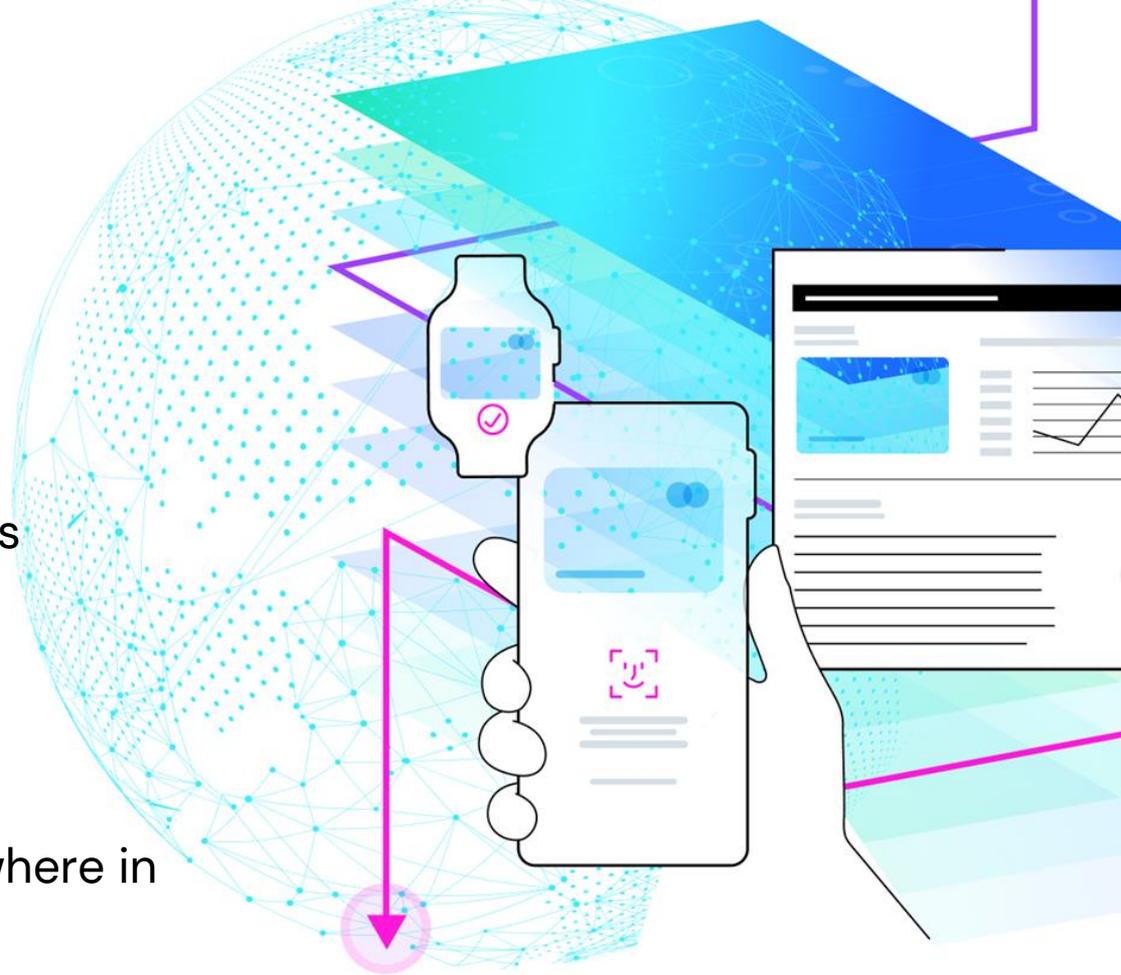
- HPE NonStop Server



PAYMENT STACK CAPABILITIES

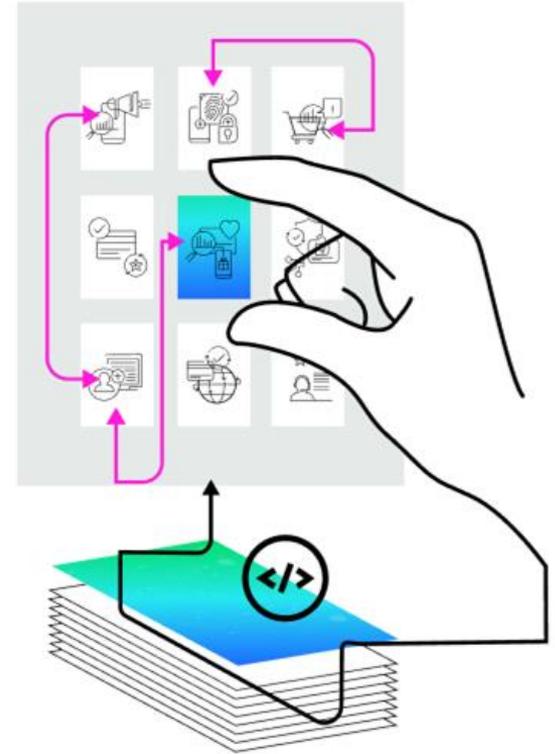
Technology that keeps you ahead

- HPE NonStop Server Clusters
- Early adopter of ServerNet Clustering
- NB56000 systems used BladeCluster products
- Applications deployed anywhere in the cluster



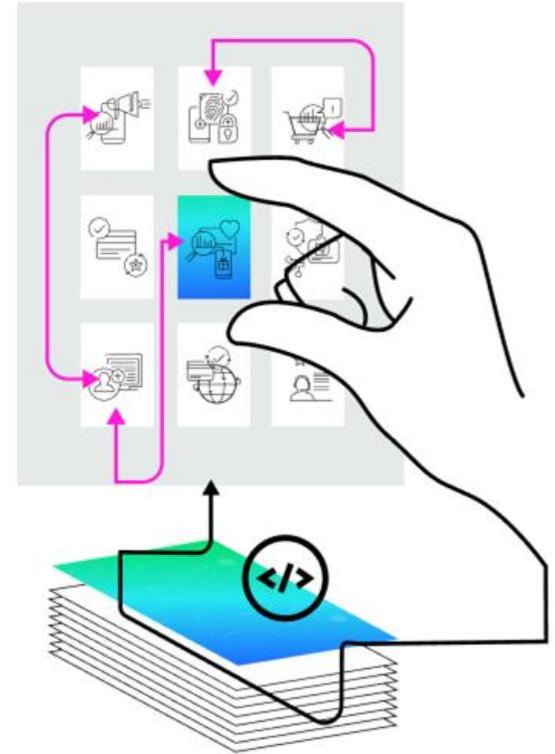
NSX Cluster Solution

- The NSX Cluster Solution is a hardware and software product which allows NonStop L-Series nodes to communicate quickly and efficiently through the InfiniBand fabrics



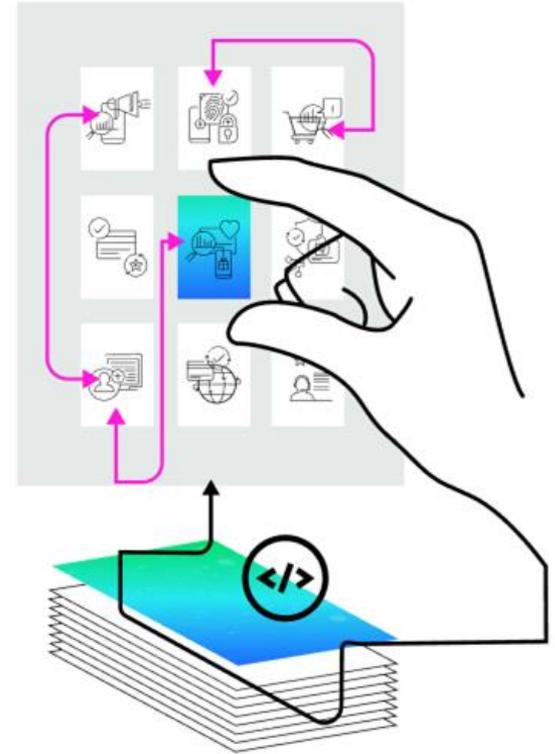
NSX Cluster Solution

- The NSX Cluster Solution is a hardware and software product which allows NonStop L-Series nodes to communicate quickly and efficiently through the InfiniBand fabrics
- Similar to earlier products:
 - BladeCluster
 - ServerNet Cluster
 - Fiber Optic Extension (FOX)



NSX Cluster Solution

- The NSX Cluster Solution is a hardware and software product which allows NonStop L-Series nodes to communicate quickly and efficiently through the InfiniBand fabrics
- Usable by NS7, NS8 and vNonStop servers
- NS7 and NS8 servers can participate in the same cluster



Why use multiple NonStop systems?

- The workload is more than one system can handle
- Better availability using active-active processing over multiple systems



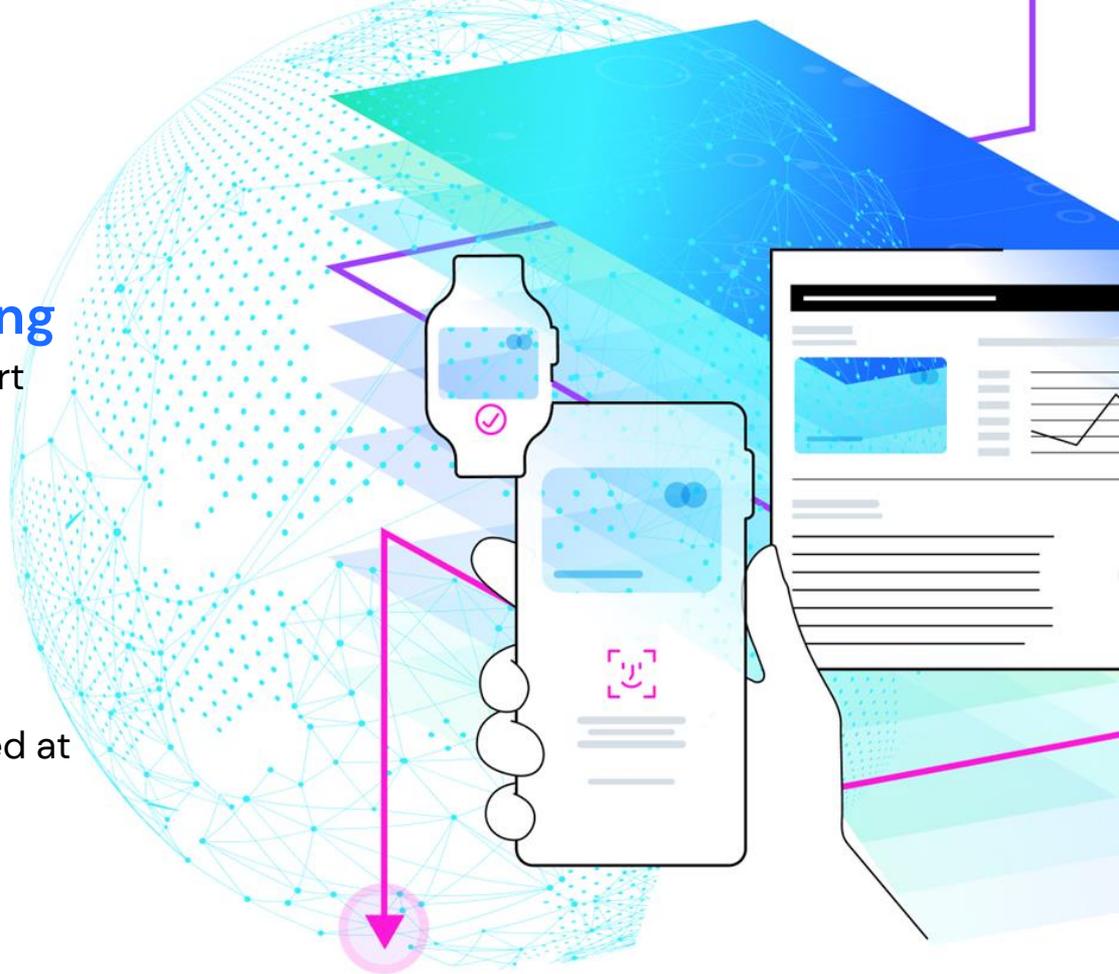
Why NSX Cluster Solution?

- The network fabric does most of the routing work
 - Significantly less CPU consumption
 - Processing distributed more evenly over all CPUs
- Better expected reliability



Cluster Migration One

- **BladeCluster Networking**
Four NB56000c nodes, 51 km apart
- **Collapse Compute**
Migrate to two NX8X4 nodes
- **Expand/IP Networking**
NSX Cluster Solution not supported at this distance



Cluster Migration One

Get detailed BladeCluster loading Baseline metrics:

- Latency - < 1 ms
- Peak throughput - 146 Mb/sec



Cluster Migration One

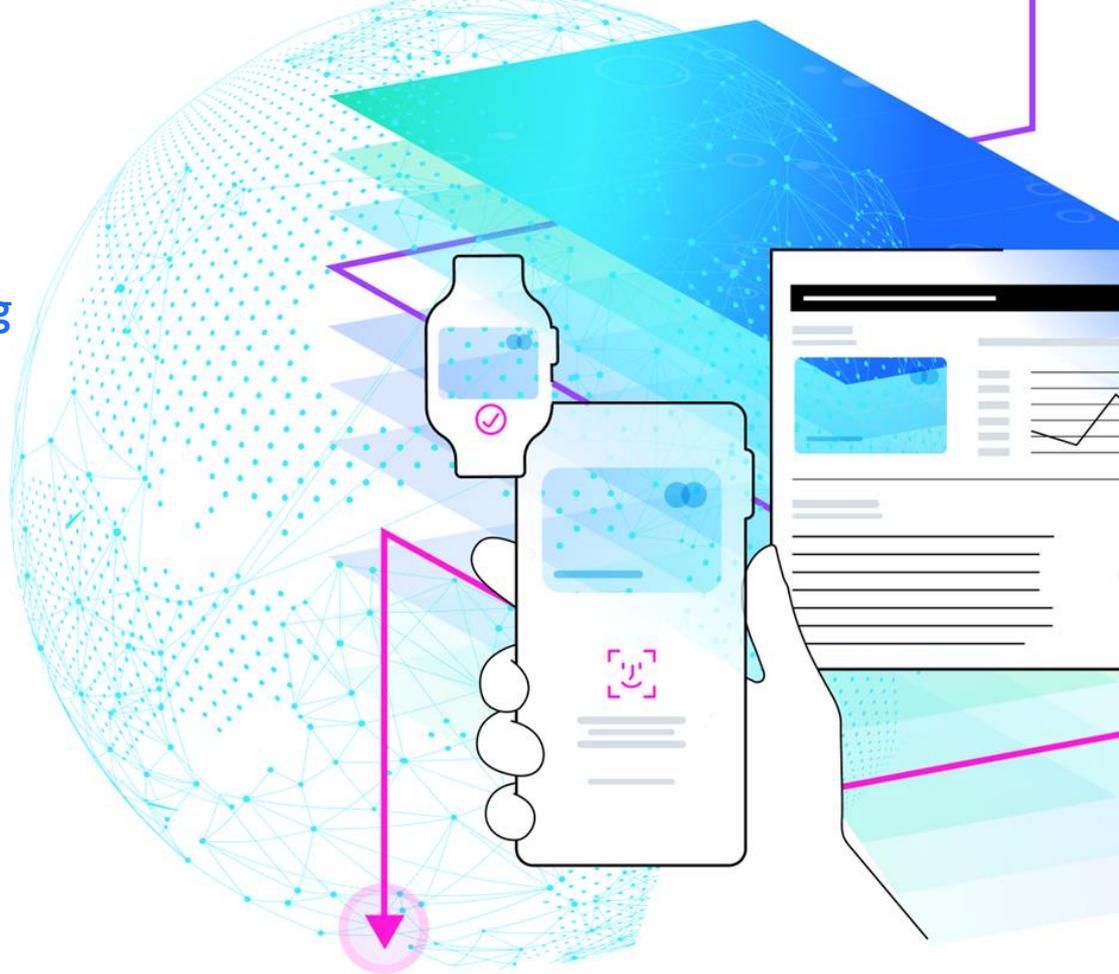
Get detailed BladeCluster loading

Baseline metrics:

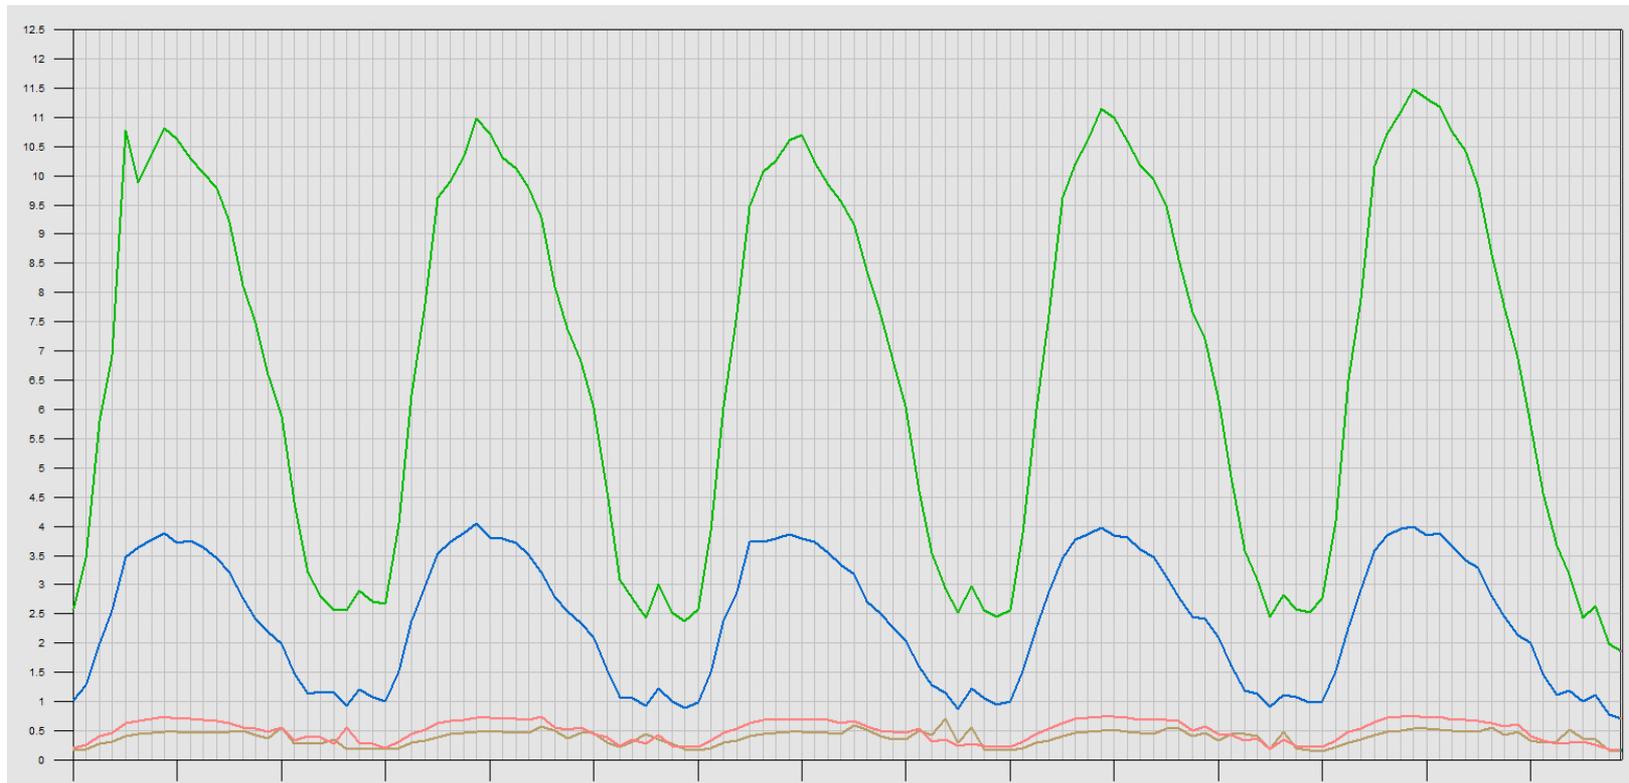
- Latency - < 1 ms
- Peak throughput - 146 Mb/sec

Expand/IP Network Architecture

- Two logical network paths
- Four logical superpath lines



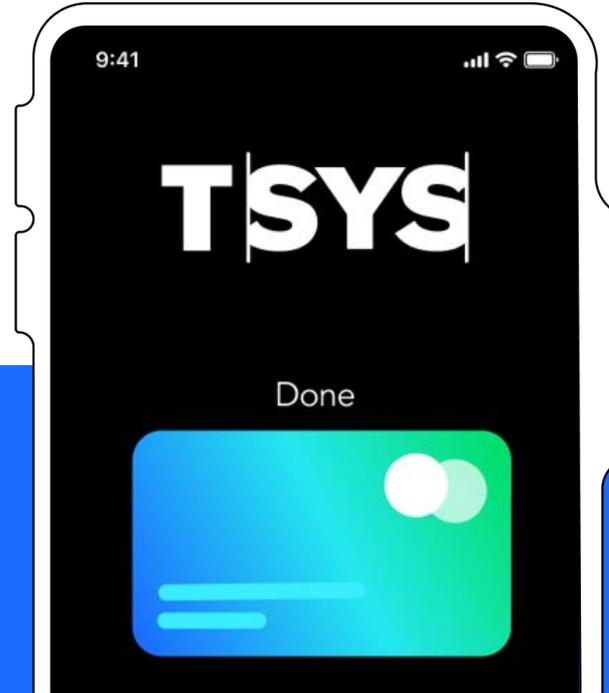
Expand/IP CPU utilization



Lessons Learned

Expand/IP Path CPU consumption

- Proportional to traffic
- Minor concern at this loading



Scale up

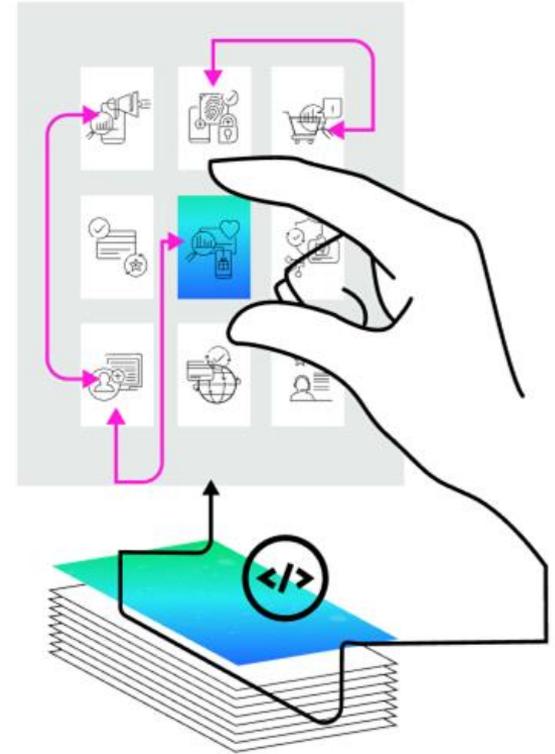
Use the analysis and processes developed with the smaller cluster and apply to a busier cluster.

Three larger nodes

- Scale up to 3 nodes.
- All nodes within 30 meters

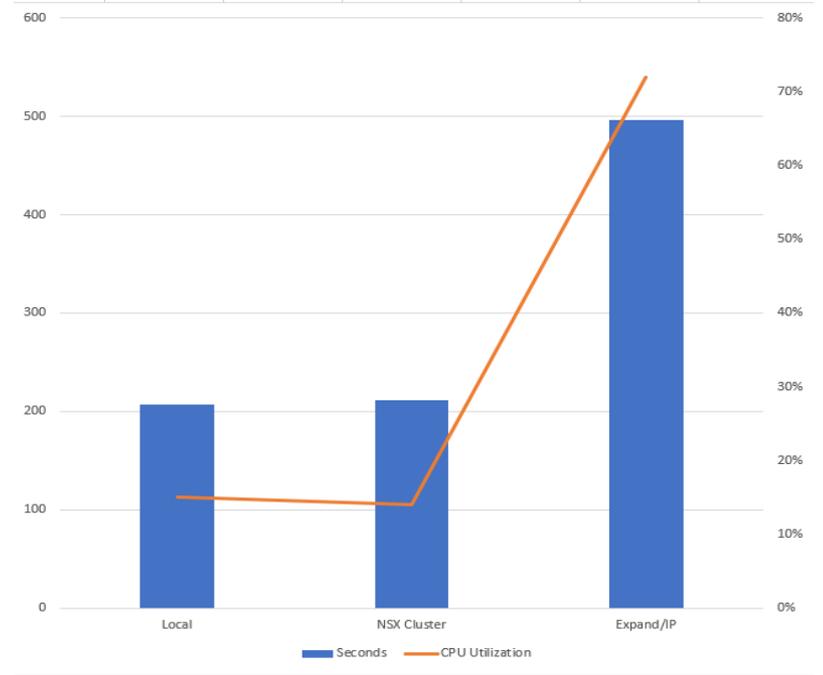
Higher peak throughput

- 1,751 Mb/sec peak throughput



Bulk Data Transfer

	Seconds	Mb/sec	CPU Busy
Local	207	1,979	15%
NSX Cluster	211	1,942	14%
Expand/IP	496	826	72%

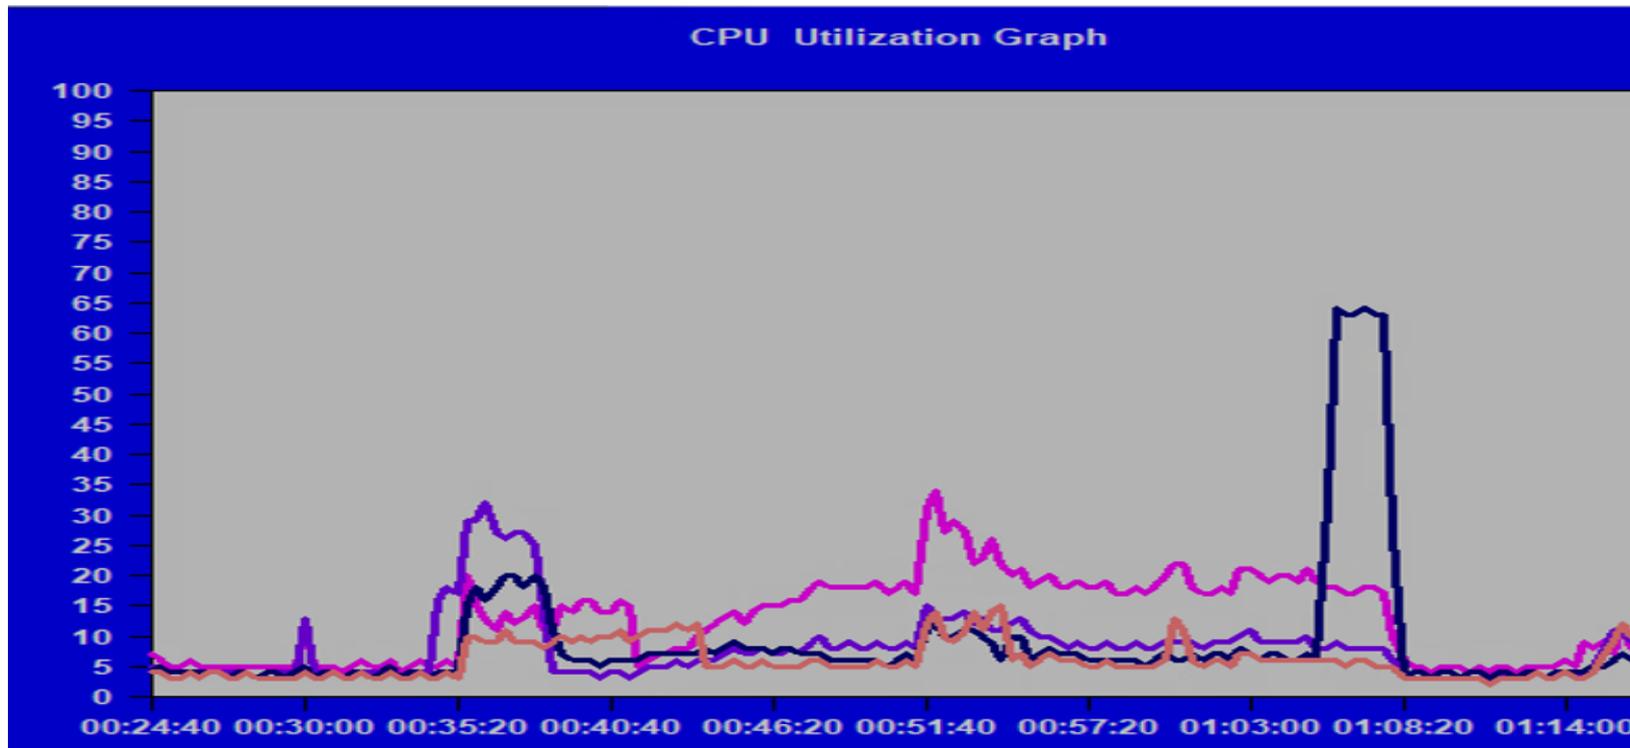


Bulk Data Transfer

NAME	BUSY %	RCV QLEN	RCV RESP	Program	PID	USER	PRI	MEM PAGE	ANC NAME	CPU NO
\$PX43	41.89	0.0	0.4806	\$SYSTEM.SYS01.LHOBJ	2,443	255,255	199	355	\$ZZWAN	2
\$:2:307	30.08	0.0	0.0000	\$SYSTEM.SYS01.TSCOMIP	2,307	255,255	255	11		2
\$REST01	7.51	0.0	0.0000	\$SYSTEM.SYS01.TSYSDP2	2,551	255,255	220	3303	\$ZZSTO	2
\$SM18	2.21	0.0	0.0000	\$VCODE.NDMOBJ.NDMSMGR	0,1497	90,255	130	67	\$NDMS	0

NAME	BUSY %	RCV QLEN	RCV RESP	Program	PID	USER	PRI	MEM PAGE	ANC NAME	CPU NO
\$REST01	13.99	0.0	0.0000	\$SYSTEM.SYS01.TSYSDP2	2,551	255,255	220	3303	\$ZZSTO	2
\$:2:293	3.14	0.0	0.0000	\$SYSTEM.SYS01.TSSTOIP	2,293	255,255	255	11		2
\$:2:323	2.99	0.0	0.0000	\$SYSTEM.SYS01.TSMMSGIP	2,323	255,255	255	11		2
\$:2:297	2.86	0.0	0.0000	\$SYSTEM.SYS01.TSSTOIP	2,297	255,255	255	11		2
\$:3:321	1.90	0.0	0.0000	\$SYSTEM.SYS01.TSMMSGIP	3,321	255,255	255	11		3
\$REST01-B	1.90	0.0	0.0000	\$SYSTEM.SYS01.TSYSDP2	3,573	255,255	220	2242	\$ZZSTO	3
\$:2:322	1.67	0.0	0.0000	\$SYSTEM.SYS01.TSMMSGIP	2,322	255,255	255	11		2

Simulated Load



Lessons Learned

NSX Cluster solution

- Very close to local throughput
- Nearly zero host CPU cost
- Even NSX pass through used less CPU than Expand/IP



Scale out

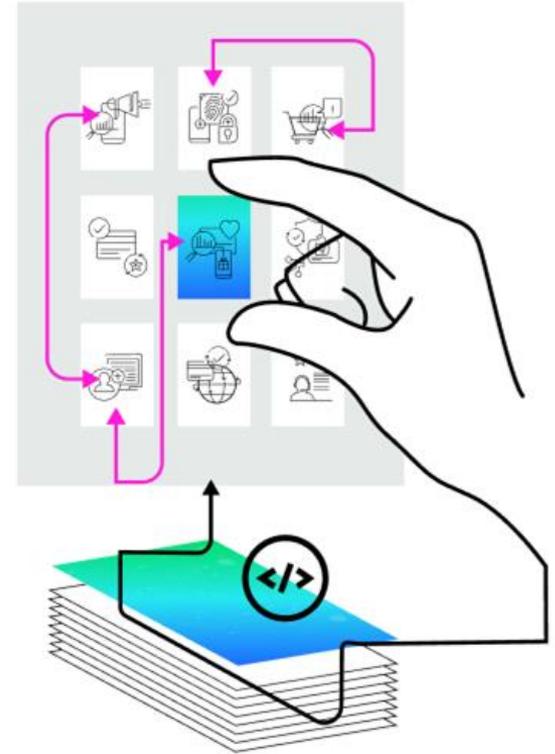
Use the analysis and processes developed with the smaller cluster and apply to a larger cluster.

Add more nodes

- Scale out to 10 nodes
- All nodes within 30 meters

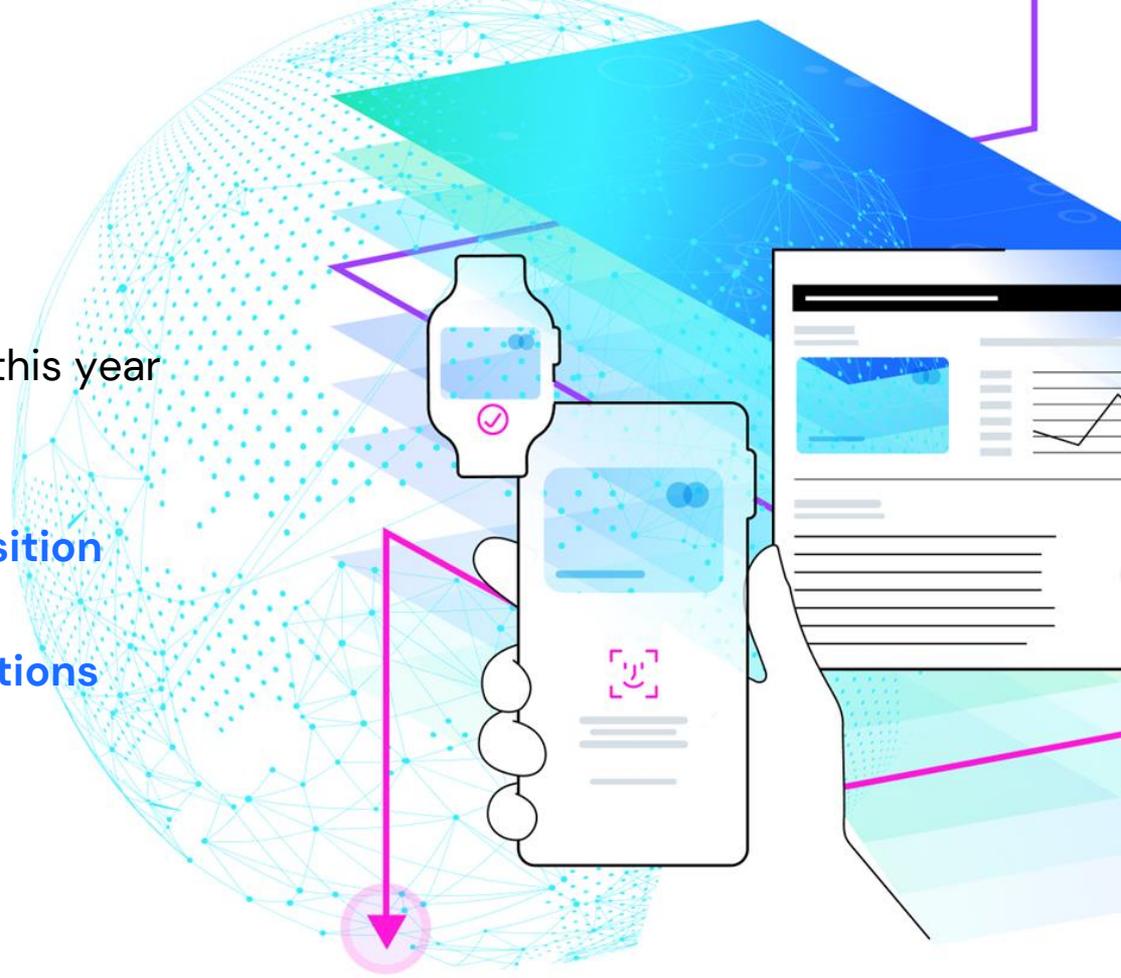
High peak throughput

- 883 Mb/sec peak throughput



Cluster Migration Three

- **Work in process –**
Four nodes to be completed this year
Six nodes planned for 2024
- **Capacity for Expand/IP transition**
- **Early results match expectations**

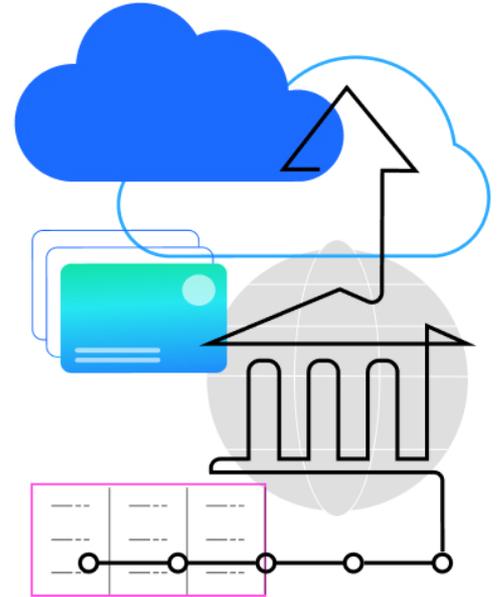


Dollars and Sense

NSX Cluster Solution

What makes it an
effective solution?

Performance justification
Much higher throughput than
alternatives



NSX Cluster Solution

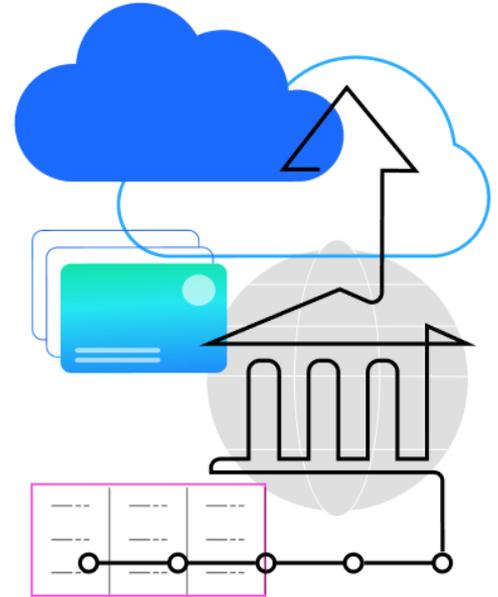
What makes it an effective solution?

Performance justification

Much higher throughput than alternatives

Better operational characteristics

Deploy applications where they become easier to manage, regardless of node boundaries



NSX Cluster Solution

What makes it a cost effective solution?

Performance justification

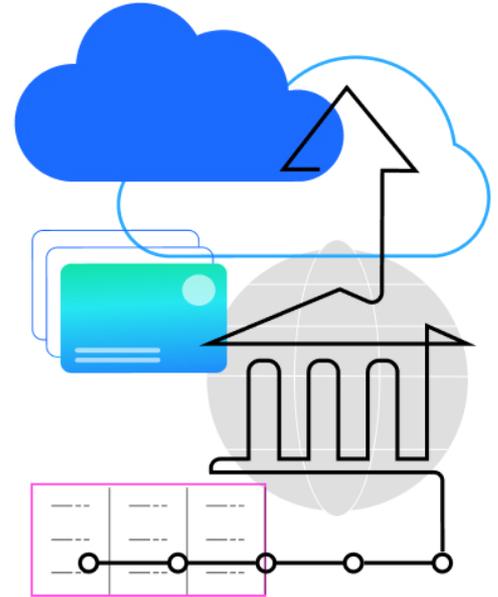
Much higher throughput than alternatives

Better operational characteristics

Deploy applications where they become easier to manage, regardless of node boundaries

Cost savings

Due to reduced CPU needs, the NSX Cluster Solution may even cost less than using Expand/IP



Summary

We've found the HPE NSX Cluster Solution to be a helpful option to handle our high-throughput and low-latency workload.



Excellent throughput

Alternative networking options are slower



Might be the lowest cost

Offload routing to the super smart NonStop network



Scalability

The bigger the workload, the more cost effective it is



Consider replication

Let database replication traffic use the NSX Cluster Solution

Questions?



Thank you

Roy L. Nicholas
RNicholas@tsys.com