

Utilization of Virtual Server Technology in Mission Operations

Larry Felton¹

CSC, Huntsville, Alabama, 35806, USA

Kimberly Lankford²

COLSA Corporation, Huntsville, Alabama, 35806, USA

R. Lee Pitts³

CSC, Huntsville, Alabama, 35812

and

Robert W. Pruitt⁴

COLSA Corporation, Huntsville, Alabama, 35806, USA

Virtualization provides the opportunity to continue to do “more with less”—more computing power with fewer physical boxes, thus reducing the overall hardware footprint, power and cooling requirements, software licenses, and their associated costs. This paper explores the tremendous advantages and any disadvantages of virtualization in all of the environments associated with software and systems development to operations flow. It includes the use and benefits of the Intelligent Platform Management Interface (IPMI) specification, and identifies lessons learned concerning hardware and network configurations. Using the Huntsville Operations Support Center (HOSC) at NASA Marshall Space Flight Center as an example, we demonstrate that deploying virtualized servers as a means of managing computing resources is applicable and beneficial to many areas of application, up to and including flight operations.

I. Introduction

The Huntsville Operations Support Center (HOSC) has evolved over time from an isolated but open system to a system which supports local and remote access by its users over diverse geographic regions. This has been accomplished to reduce cost and provide the highest levels of user support. In the last 10 years the HOSC has been converted from a client server system supporting the Shuttle and Spacelab on isolated proprietary systems, to one that now has fully enabled Internet Protocol (IP) applications.

Specifically, the HOSC hosts the ground system for the U.S. portion of the International Space Station (ISS), processes and reduces Shuttle data, is the authoritative source for Ares Flight data, and supports numerous small operational projects such as the Fast Affordable Science and Technology Satellite-Huntsville01 (FASTSat-HSV01). The Chandra X-Ray Observatory utilizes a version of the HOSC ground system. As the HOSC has expanded its role as a NASA operations center, it has embraced open systems solutions on commodity platforms and now hosts most applications on LINUX or Microsoft Windows environments.

At the same time, the performance of Complex Instruction Set Computer (CISC) servers has dramatically increased with the addition of 64 bit processors, additional memory, higher speed clocks, multiple cores, and many Reduced Instruction Set Computer (RISC)-like features. As a result, one of the newest technologies the HOSC is deploying in operations and support is virtualization.

This paper proposes that the use of virtual server technology is not only cost effective but viable in space operations. The HOSC ground system complex currently utilizes virtual server technology and is deploying it in various areas in direct and indirect operations. This paper specifically addresses the direct deployment in operations.

¹ Systems Engineer, 310 The Bridge Street.

² HOSC Operations and Maintenance Manager, MSFC/Mission Operations Laboratory (MOL).

³ HOSC Chief Engineer, MSFC/MOL, AIAA Senior Member.

⁴ Operations Engineer, MSFC/MOL.

II. Identifying an Opportunity

Since the inception of the Enhanced HOSC System (EHS) to support the ISS, the HOSC has moved away from centralized mainframe computing to a client-server model. In 2000, the EHS supporting the ISS was deployed. The initial deployment was across SGI, SUN, and DEC equipment of multiple generations with the occasional IBM thrown in. Since then, the EHS footprint and cost have been successively reduced through modernization, commodity hardware, and a Storage Area Network (SAN) implementation. However, to a large extent the initial deployment paradigm is still in place; one logical set of applications to a physical server, one physical workstation per user view of the system.

Beginning in 2003, the HOSC began migrating to LINUX on CISC for its mission application servers. This was an extensive retrofit activity touching nearly every aspect of the HOSC. It resolved a number of issues and provided benefits such as:

- Removing obsolete capabilities
- Consolidating similar functions
- Inserting new technologies such as a Storage Area Network
- Changing the overall server philosophy with commodity priced components
- Isolating discrete and obsolete devices
- Enhancing security

Four initiatives were completed that resulted in major cost savings and performance increases. As a result the HOSC is able to do more at lower cost by expanding its server base. This includes supporting more activities and users as well as providing new and richer capabilities at lower costs.

The HOSC architecture is composed of a number of Virtual Local Area Networks (VLANs) that support the management of security and user access. Operations is spread across these VLANs based on Operations/Simulation/Independent Validation and Verification (IV&V) functions and whether a user of the services is local or remote. Core systems are more highly protected and deeply embedded within the architecture.

From a network standpoint, the HOSC comprises an inner and outer domain with devices allocated to subnets based on usage. Figure 1 depicts the various HOSC network tiers.

In the inner domain network, approximately ten (10) subnets exist. Specifically, and not including those networks supporting Shuttle Data Reduction (STS DR) or the ISS data ingest network Payload Data Services System (PDSS), the subnets are:

1. Payload Ops – Core Command and Control capability
2. Dev / Val – IV&V for Payload OPS
3. MSS – Mission Support Services for infrastructure items
4. EPC – Internal clients using Enhanced PC to EHS systems
5. PVT – Private LAN hosting client services for local clients
6. ePVT – external Private LAN hosting client services for remote clients

Several other subnets are inaccessible to most users including remotes. These networks support systems and network management and the management of realtime data access. Specifically, the term subnets is used because access to the HOSC services is by subnet as specified by the Interface Control Documents (ICD).

There are a number of generic services which are used by many HOSC activities and are found on several subnets. First, there are a number of file servers for services such as Near Real Time (NRT) data access, users' mission products, and other items unique to the processing of mission data. These are hosted almost exclusively on an Archive LAN. Additionally, other servers are hosted on a mission support VLAN and other subnets. These servers include Lightweight Directory Access Protocol (LDAP), Internal Domain Name System (DNS), Database build servers, operating system management servers, and test servers. Many of these servers are invisible to external interfaces.

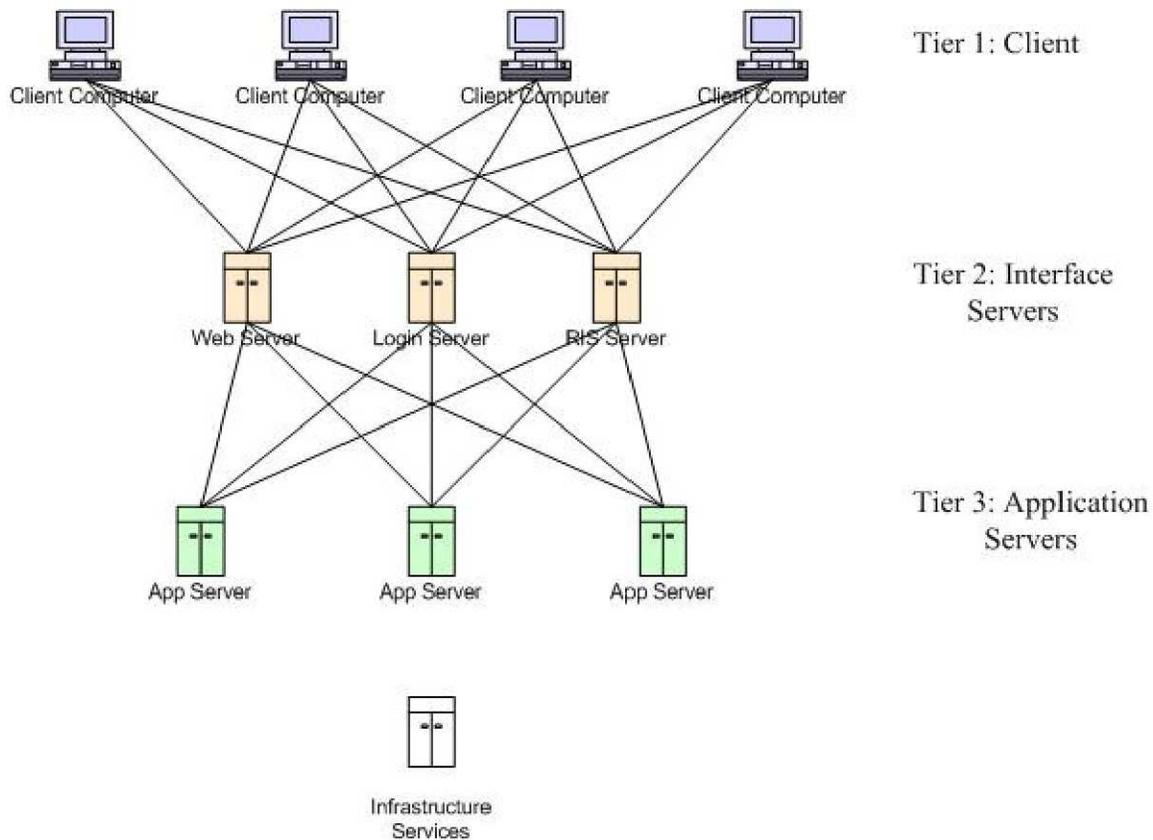


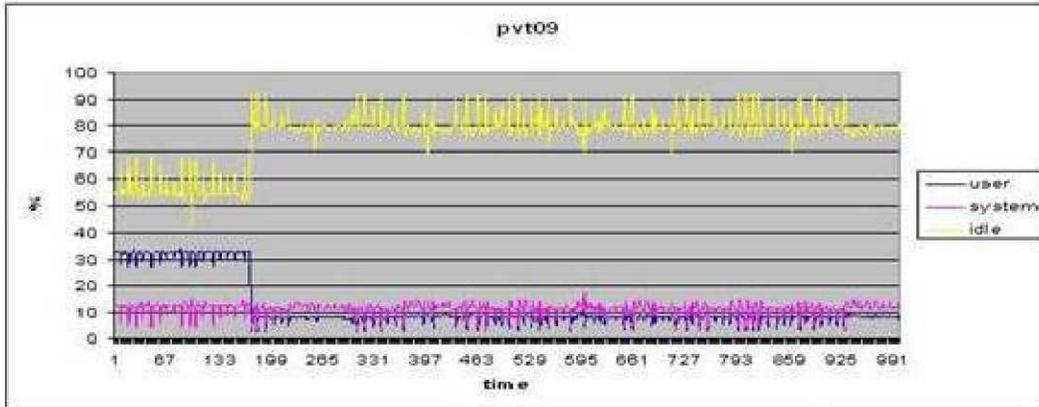
Figure 1. HOSC Network Tiers.

Finally, external access to HOSC systems is via three (3) subnets:

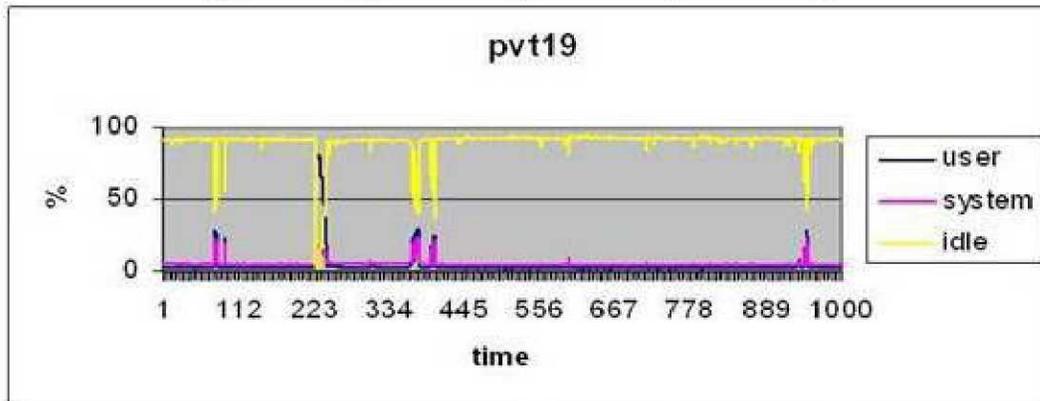
- PVT, which hosts Ground Support Equipment and Dataset packets, NRT playbacks, and NRT-related File Transfer Protocols.
- ePVT, which hosts external user access to include Enhanced PC, X-windows, and programmatic commanding.
- HOSC DMZ, a firewall "Demilitarized Zone" that hosts the HOSC web services.

Systems on these three subnets are the initial targets for virtualization. The Deploying Virtual Servers section provides additional information concerning the primary reasons that these servers are candidates for virtualization. Each of these servers supports the client tier.

While examining performance metrics it was noticed that the new servers were not taxed in their resource utilization as shown in Fig. 2. In fact, the servers had utilization well below 50 percent. In the current configuration and layout, they primarily provide "peak" capacity, redundancy, and fault tolerance at bargain prices.



| | user % | System % | idle % |
|---------|--------|----------|--------|
| average | 12.02 | 11.52 | 76.10 |
| Max | 33.93 | 17.24 | 92.7 |
| Min | 2.84 | 4.1 | 42.83 |



| | user % | System % | idle % |
|---------|--------|----------|--------|
| average | 4.20 | 5.28 | 90.05 |
| Max | 80.91 | 23.52 | 94.16 |
| Min | 1.9 | 3.69 | 1.57 |

Figure 2. PVT Server Performance Examples.

Specifically, the servers are standard dual Intel processor servers (< 3 Gigahertz (GHz)) with 4-8 gigabytes of memory. In their operational configuration the servers support user applications such as telemetry and command and control. Other characteristics include:

- Redundant power supplies sourced to redundant power sources (circuits) which can individually supply a server.
- Single Redundant Array of Disks (RAID) 10 disk configuration.
- Single Network Interface Controller (NIC) per VLAN supported for up to four (4) NICs per server.
- Physical separation by 30+ meters to limit failures related to environmental coupling, e.g., air flow, fire suppression shower heads, loss of a circuit panel.
- Red Hat™ 4.x LINUX operating system and various COTS products and locally developed custom software.

The HOSC has a large quantity of these servers, each of which is configured slightly differently to support operations, simulation, and test. Even though the basic configuration is the same, each has unique aspects based on its usage.

Additionally, the overall system architecture has several failure points that are untenable when supporting virtualized systems, i.e., multiple instances on a single physical instance. A higher level of server availability is required to be implemented through the use of licensing, robust hardware configurations, and operational configuration. Areas of concern are:

1. Critical networks' access (single NIC and switch port)
2. RAID configuration that would be intolerant of virtual servers
3. Memory bandwidth
4. Method of utilization (Operations Concept)
5. Physical location of servers

III. Deploying Virtual Servers

Deploying virtual servers requires an understanding of the current server operations concept and their configuration. In the HOSC, operational servers are assigned to an activity or mission. As depicted in Fig. 3, ISS Mission Operations has a complete and fully redundant set of resources while the simulation string using the same software build may have a reduced non-redundant set. Another complete set of resources is available as "transition resources." These transition resources host the next software build that will be used operationally.

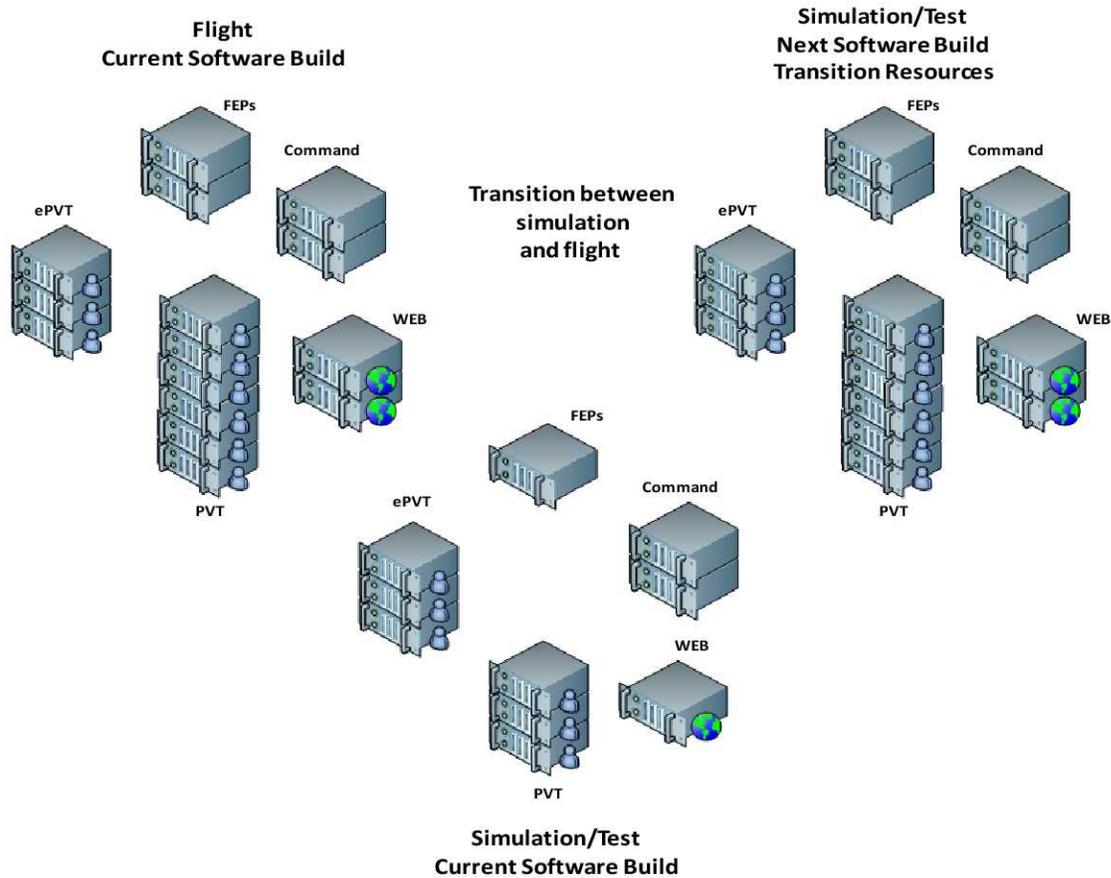


Figure 3. Current HOSC Server Allocation.

In Fig. 3, the PVT, ePVT, and Web servers (Tier 2 resources in Fig. 1), are prime candidates for virtualization. This is primarily because they are a server pool shared between mission, simulation, and test activities. They represent a large portion of the hardware required for mission support; thus they will reap the biggest server savings with virtualization. Also, they support the user interfaces--versus infrastructure interfaces--and a diverse set of configurations are available for inclusion. The failure of a single redundant platform will not adversely affect ongoing operations or necessarily a single mission center.

A. Proving the Concept

Several platform characteristics were evaluated to ensure heterogeneous environments are not deleterious to operations and are transparent to operations. Specifically the areas of concern are:

- Platform status, both logical and physical
- Isolation of logical platform performance characteristics
- The ability of on-shift personnel to manage and respond to anomalies without engineering support
- Right platform for the right mission; do not waste or constrict resources

The necessary platform characteristics that must be evaluated and possibly bolstered are:

- Onboard disk storage and its survivability and separability (RAID usage)
- Processor characteristics and number
- Onboard RAM and organization
- Network accessibility and availability of critical networks to real time
- Virtualization characteristics ("hardened" and inviolate)
- Mission boundaries affected by virtualized platforms
- Physical server configuration

A candidate configuration was developed and a series of tests were devised to evaluate the use of virtualized servers in operations at the HOSC. Two servers were selected to implement virtualization in the test environment. They were selected to provide three virtual servers; each approximates the current configuration utilized in simulation and operations.

The servers were configured in a rough approximation of the current operational systems. Figure 4 compares an operations server to the proposed server configuration. Though not exact, similar configurations permit more direct comparison of behavior.

| Component | Operations | Physical | Virtualized |
|-----------------|-----------------------|----------------------------|--------------------|
| Motherboard | SM X5DL8 | SM X7DB3 | N/A |
| Processor | 2x Intel® Xeon® CPU | 2x Intel® Xeon® X5355 CPU | 2x |
| Processor speed | 3.06GHz Cache: 512 KB | 2.66 GHz Cache: 4 MB | N/A |
| Platform cores | 1x2 | 2x4 (2) | 2 |
| FSB | 533 MHz | 1333 MHz | N/A |
| Memory | 4 GB | 16 (4) GB | 4 GB |
| RAID | RAID 10 (37) | RAID 10 (10) | RAID 10 (33) |
| Mission network | single, 100 Mbps | 1000 Mbps channel bond (1) | Virtualized |
| Archive network | single, 100 Mbps | single, 1 Gbps | Virtualized |
| Other network | N/A | single, 1000/100 Mbps | Virtualized |
| Operating Sys | LINUX RH 4.6 | LINUX RH 5.4 (HOST) | LINUX RH 4.6 Guest |

Figure 4. Current Operations Server to Proposed Server Comparison.

From a partition standpoint, the hosted virtual servers are the equivalent of one physical server in operations with identical connectivity on Gigabit interfaces that are shared with all other virtual servers on a single physical platform. As depicted in Fig. 5, the mission side is annotated as the PVT VLAN. This is the interface that supports mission operations. PVT servers are utilized by local users to interact with client server applications. The PVT is composed of two (2) channel bonded mode1 NICs across separate Cisco switches. An ARCHive VLAN, Intelligent Platform Management Interface (IPMI) VLAN, and a Payload Data Distribution (PDD) VLAN are included for completeness.

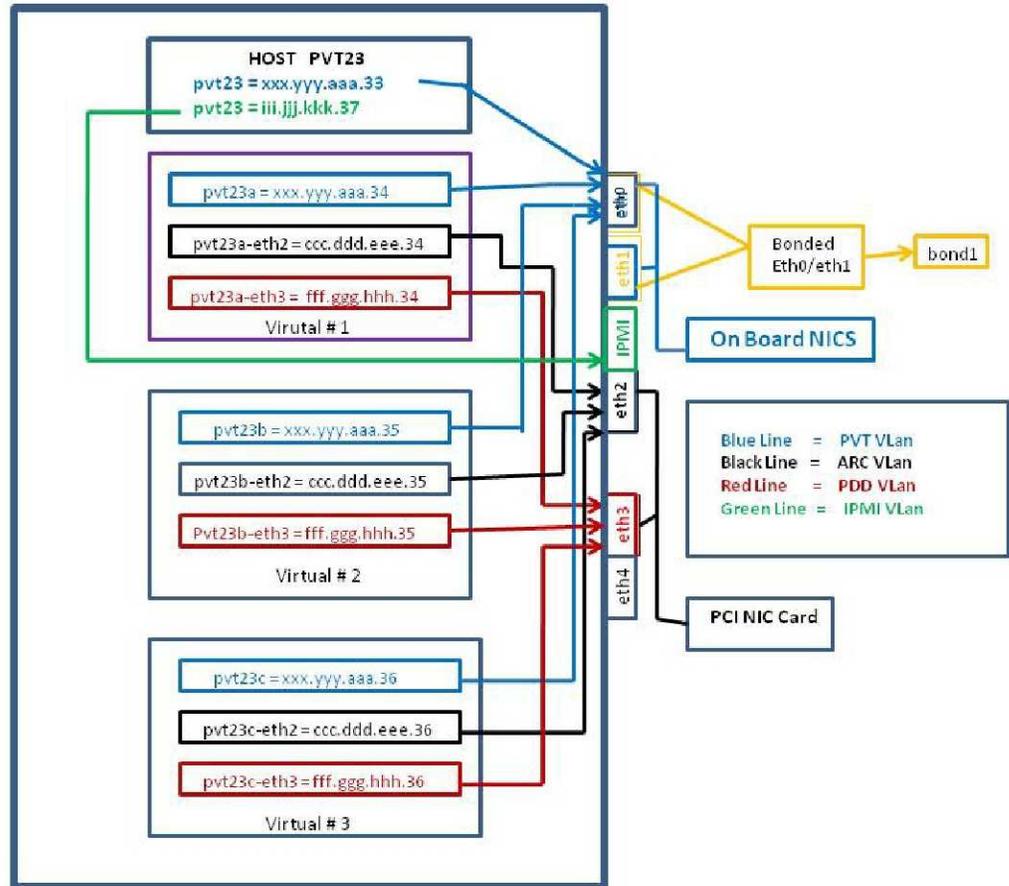


Figure 5. Virtual Server Network Layout.

Figure 6 illustrates the server layout and allocation. Each virtual server provides the equivalent capabilities of a current physical operational server. The test team loaded each virtual with the HOSC's test load. The load was executed for relatively long durations (in excess of a week). It simulated the maximum load of a "power user" in each partition. Once the servers appeared stable and processing was commensurate with normal operations, destructive testing was conducted.

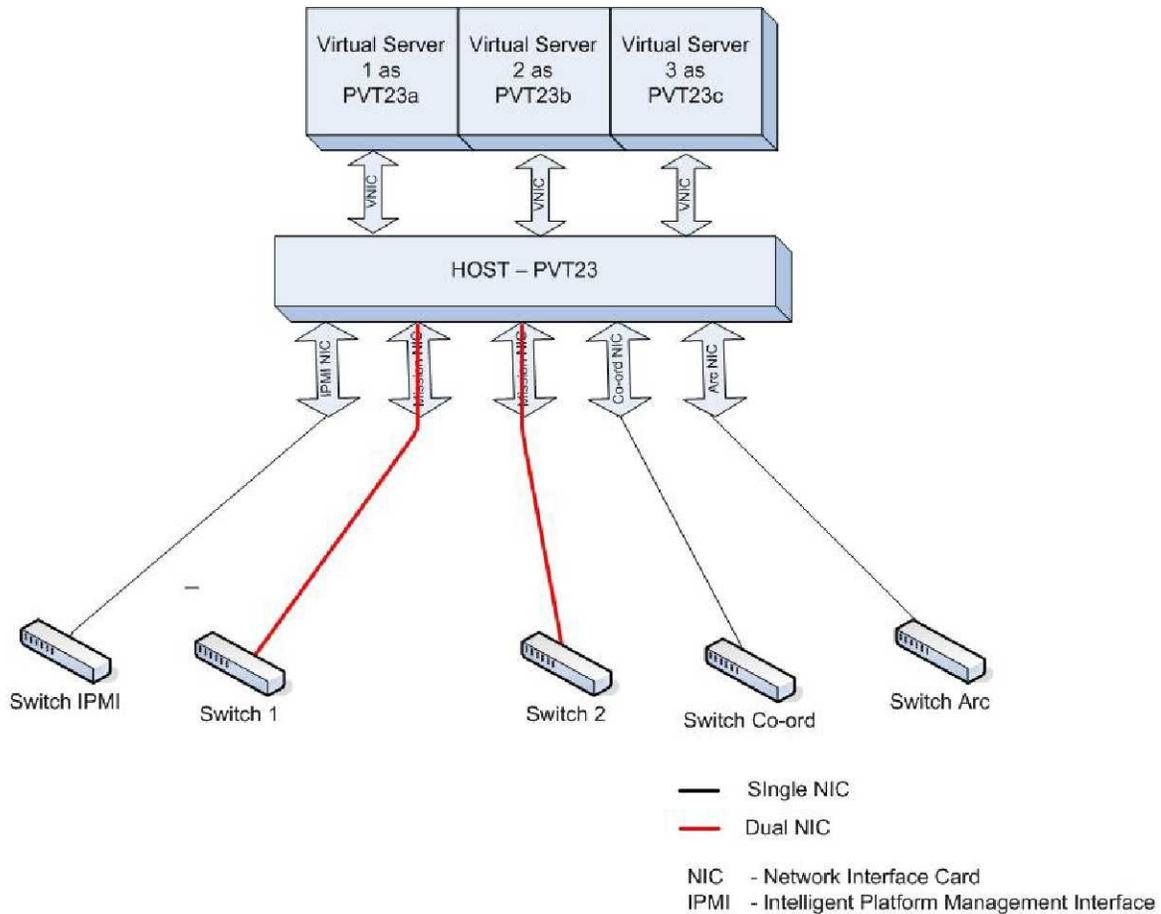


Figure 6. Physical to Virtual Server Layout.

The destructive testing included maximum memory loading and CPU loading. This was accomplished through the use of various development toolsets. Virtual servers PVT23b and PVT23c were repeatedly crashed to evaluate the effects on PVT23a. During all cases, either PVT23b or PVT23c, or both would crash. A virtual server reboot was required to restore the virtuals. The host and the virtual PVT23a did not appear to be impacted at any time after multiple attempts. In fact, testers found the virtuals more stable than a traditional server. This is attributed to the use of virtual devices.

Final testing concluded with network loss and the server response, particularly the bonded mission NICs. The original configuration for testing was Mode0 or “Round Robin”. It was found that this provided inconsistent results during normal and failover operations. Processing at times failed to be contiguous and users were required to restart processing during failure. Mode0 was also not recommended by the Operating System provider, Red Hat™. Model1 was investigated as the next viable option and accepted as a reasonable operational mode.

B. Defining the New Architecture

When investigating virtual servers, the philosophy of server allocation was brought into question. The HOSC was designed with the philosophy of a basic infrastructure shared by all with dedicated and restricted mission resources. The mission resources were servers (which may be pooled) and interfaces such as serial ports or LAN segments. Therefore, the initial virtualization model was constructed along those lines and Fig. 7 illustrates that philosophy. Operational resources were declared off-limits to simulation and test users. Operational resources represent a fully redundant set of hardware. In the current HOSC ISS model, approximately 15 servers are utilized of which 11 are viewed as candidates for virtualization. Of the remaining 4, two more are considered viable, but out of scope *at this time* due to security issues and complexity. The other two servers host special purpose and non-standard hardware which may be difficult to integrate into a virtualized system.

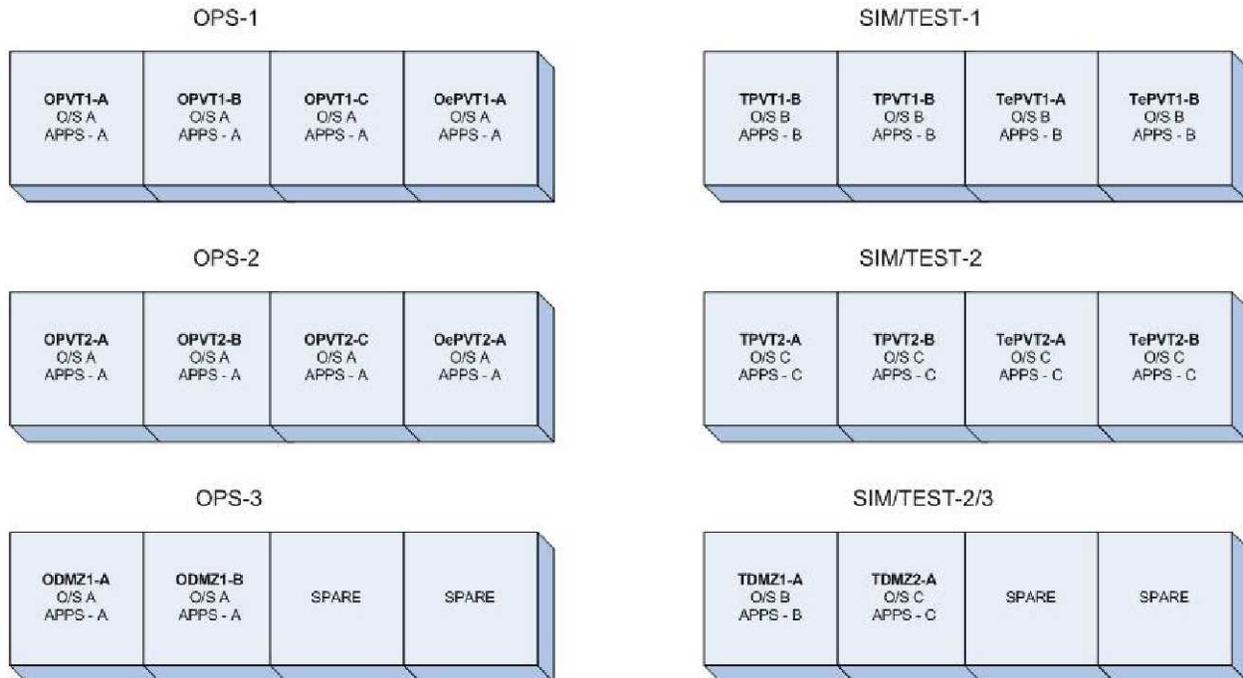


Figure 7. Initial HOSC Virtualization Model.

When the 11 servers are virtualized, they can be spread across three physical servers with in-line sparing. The operational string virtual servers represent a monoculture of a single release level of operating system (O/S A) and application code (APPS-A) release, irrespective of the underlying host operating system. Likewise, two tests or simulations can be spread across three servers. The test environment need not be a monoculture resulting in multiple operating systems (B or C) and application instances (B or C) on the virtual servers. The utilization ratio is on the order of 3 to 11 or nearly a 75 percent reduction. However, in the real world, the loss of a single physical server might cripple operations. As an operations facility, it is desirable to have physical, electrical, and communication diversity to limit the potential for a serendipitous event, e.g. inadvertent power down of a device. It was also found during testing that the destruction or corruption of the virtuals seldom--if ever--had a deleterious effect on the host. Another factor that affected the final outcome was that virtuals, once an appropriate image is built, can be managed to a detailed level, including location.

Therefore, it was proposed that operational strings be spread across a larger pool of hardware than was previously envisioned. This allowed a higher degree of diversity with less risk of a catastrophic physical server failure affecting operations. An example is shown in Fig. 8. Actual server numbers go up, but diversity is maintained and no single failure will cripple operations or test.

Figure 8 also illustrates that virtual servers can be at any operating system or application release level. This allows sparing on the individual platforms for any of the activities conducted. The ratio of utilization is 9:22 providing a 59 percent reduction in platforms with a higher degree of diversity and sparing.

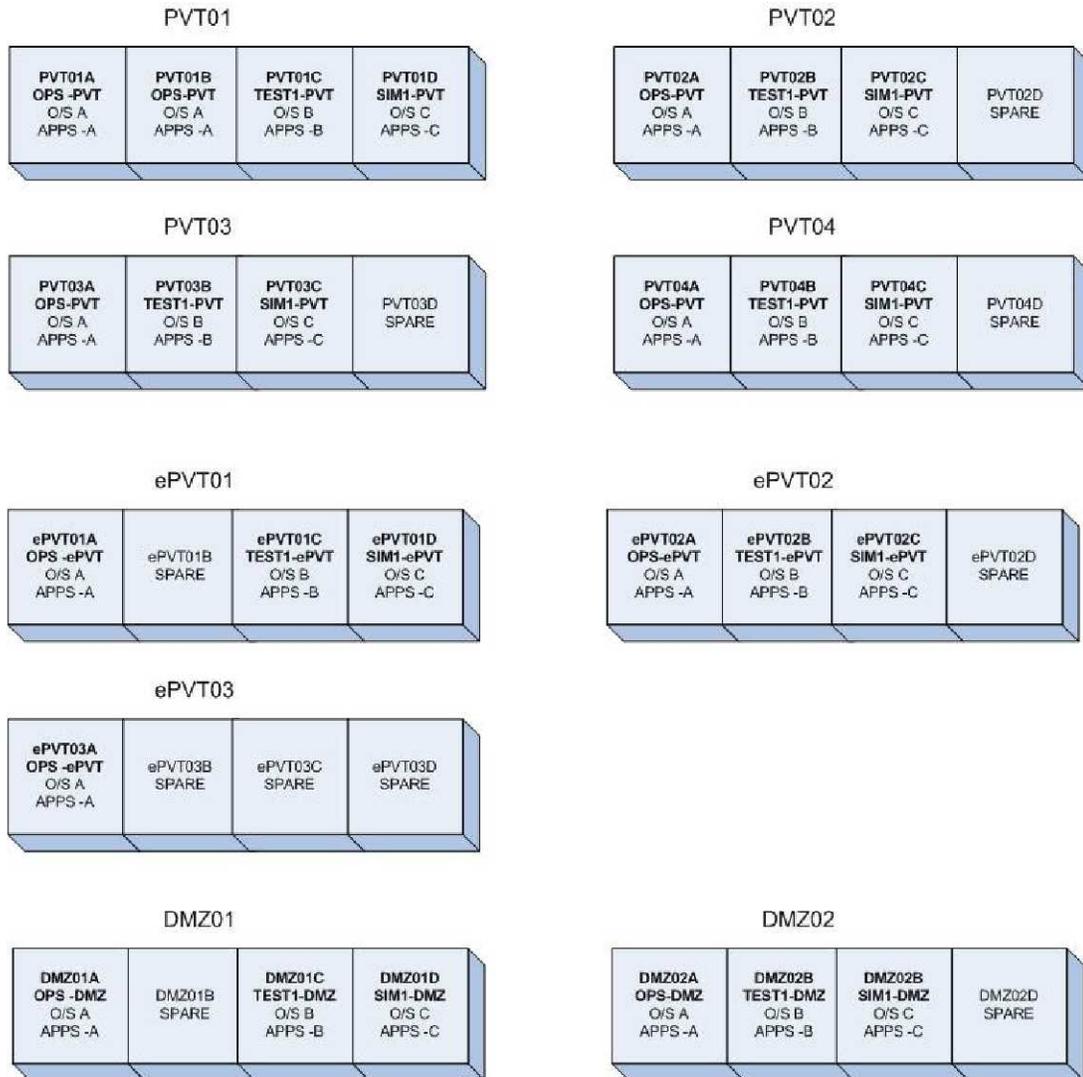


Figure 8. Diversified Usage.

The HOSC next step was to deploy the system into IV&V, the HOSC integration and final phase of testing prior to use in simulations both internally and externally. The two virtuals, PVT21 and PVT23, were totally integrated into the test environment to ensure the HOSC software subsystems responded to software reconfigurations, software installations and subsystem failover testing. This testing resulted in no software changes. The testing concluded with one virtual server being reallocated to the simulation resource pool for use by the ISS internal and external users. The server has been operational for over two months and continues to be a very stable platform. Steps are underway to configure the web server virtuals (DMZ servers) for testing. Additionally, a load balance concept is being evaluated for a CPU intensive application, the Near Real-Time retrieval application. Positive results are expected from these next series of virtual server opportunities.

IV. Conclusion

Virtualization within the HOSC facility has proved that more can be done with less. With the capability to host multiple operating systems on a single platform, the HOSC is realizing dramatic hardware reductions and cost savings in all environments:

- Development was reduced from 35 servers to 10 servers hosting virtuals, also reducing the Redhat™ software license requirements.
- Test, simulation and operational, with continued migration, will realize a 59 percent savings when all virtuals are deployed in late 2010.
- Developmental reductions will continue, allowing further cost savings in the project lifecycle.

No disruptions related to virtualization in operations are expected as a result of careful planning, testing, and implementation.

Figure 9, HOSC Virtualization Realized, illustrates the potential configuration following full virtual deployment.

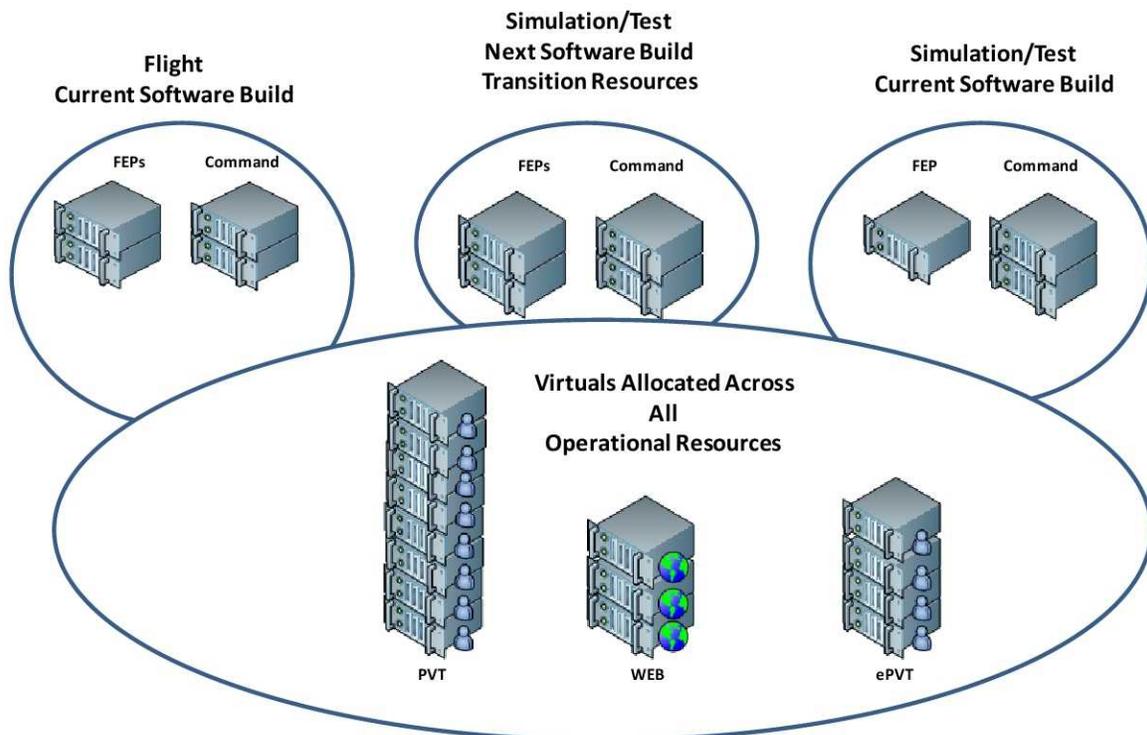
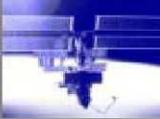


Figure 9. Virtualization Realized.

Acknowledgments

This work was supported by NASA Marshall Space Flight Center Contract NNMO4AAO7C, Huntsville Operations Support Center.



Utilization of Virtual Server Technology in Mission Operations

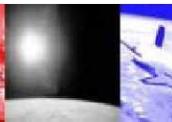
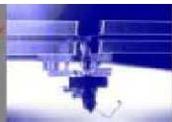


Kim Lankford
NASA Marshall Space Flight Center
April 28, 2010



Agenda

- Introduction
- Identifying an Opportunity
- Deploying Virtual Servers
- Conclusion





- Huntsville Operations Support Center

INTRODUCTION



Huntsville Operations Support Center (HOSC)



- Supports multiple programs and missions
 - Chandra
 - Initially developed, integrated, and deployed ground system
 - International Space Station
 - Host ground systems for all payloads in the US portion
 - Space Transportation System
 - Data processing and reduction
 - Ares I-X
 - Long-term archive and retrieval
 - Authoritative data source
 - Ares Test Facilities
 - Long-term archive and retrieval
 - Authoritative data source
 - Fast Affordable Science and Technology Satellite – Huntsville01 (FASTSat-HSV01)
 - Real-time and post processing



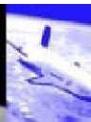
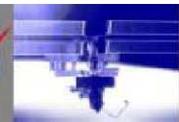
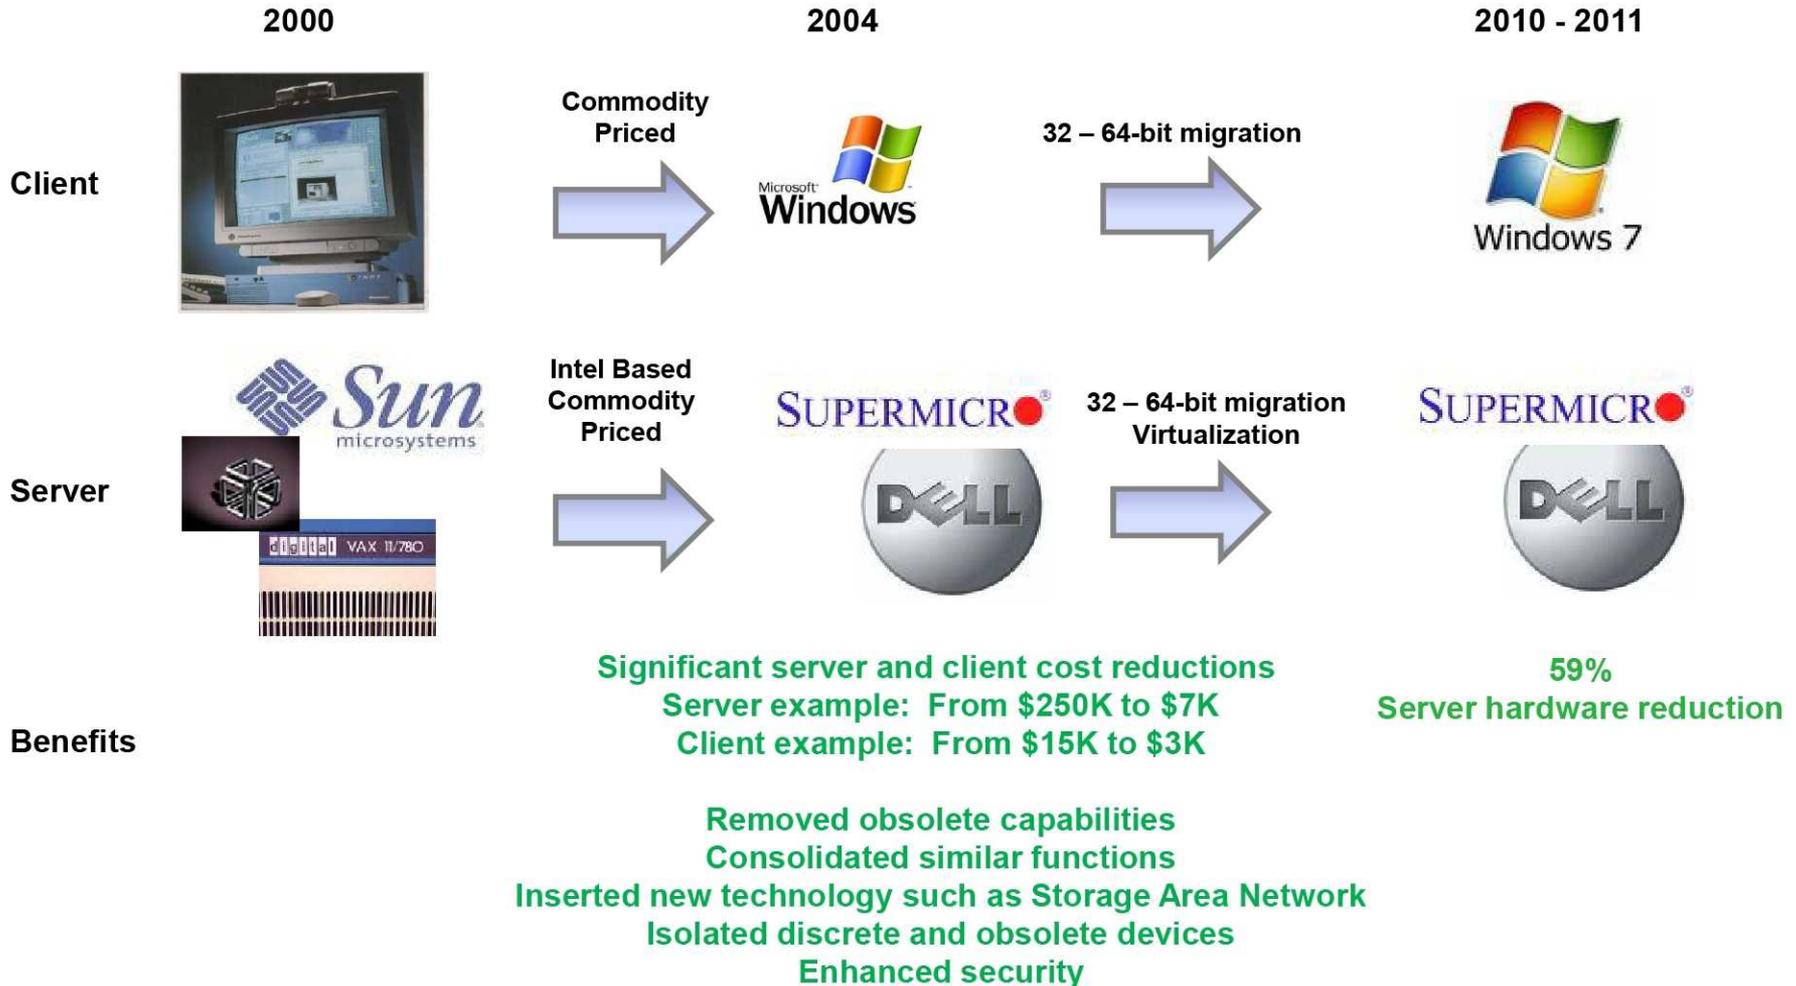


- HOSC Evolution
- HOSC Network Tiers
- Server Performance

IDENTIFYING AN OPPORTUNITY

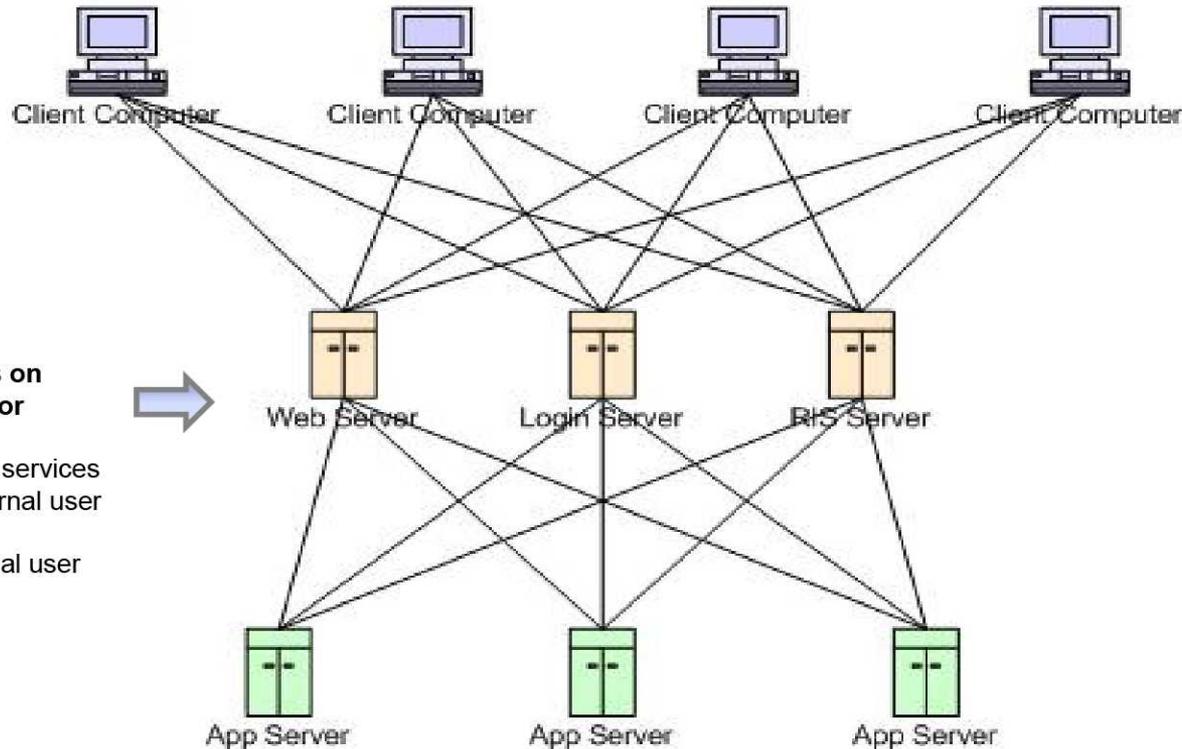


HOSC Evolution





HOSC Network Tiers



Tier 1: Client

Tier 2: Interface Servers

Tier 3: Application Servers

Three server types on subnets targeted for virtualization

- **Web** - Hosts web services
- **Login** - Hosts internal user application access
- **RIS** - Hosts external user application access

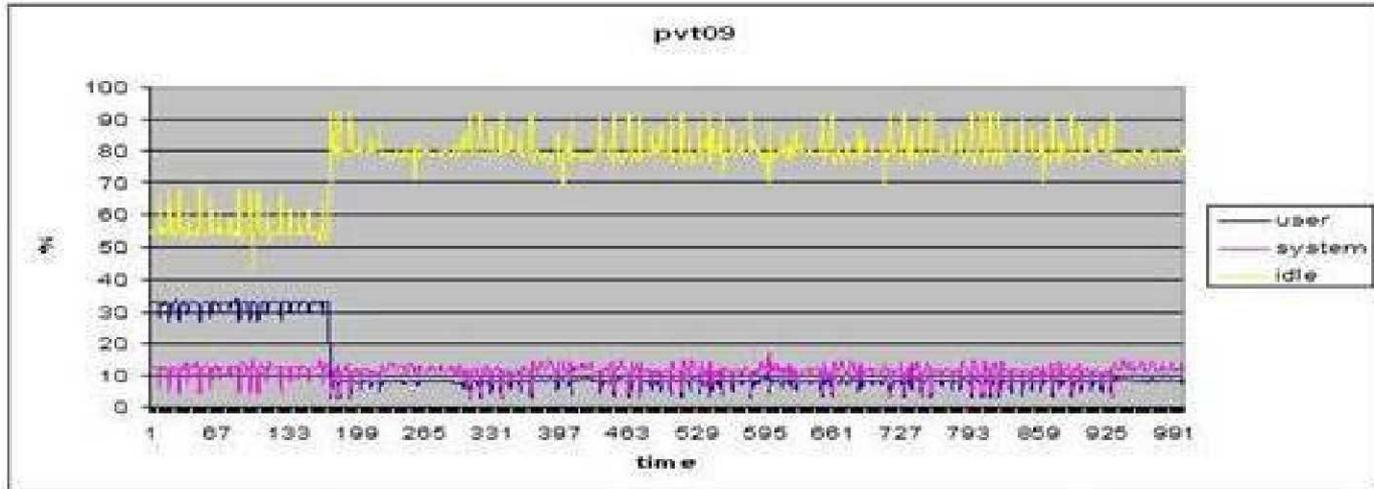
Reasons these were chosen as initial virtualization opportunities

- ✓ Represent a large portion of the hardware required for mission support; thus these will reap the biggest server savings
- ✓ No special hardware configuration required
- ✓ Support the user interfaces – versus infrastructure interfaces
- ✓ Minimum hardware resource requirements



Server Performance

- Provide “peak” capacity, redundancy, and fault tolerance at bargain prices



| | user % | System % | idle % |
|---------|--------|----------|--------|
| average | 12.02 | 11.52 | 76.10 |
| Max | 33.93 | 17.24 | 92.7 |
| Min | 2.84 | 4.1 | 42.83 |



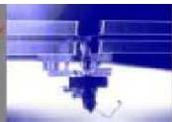
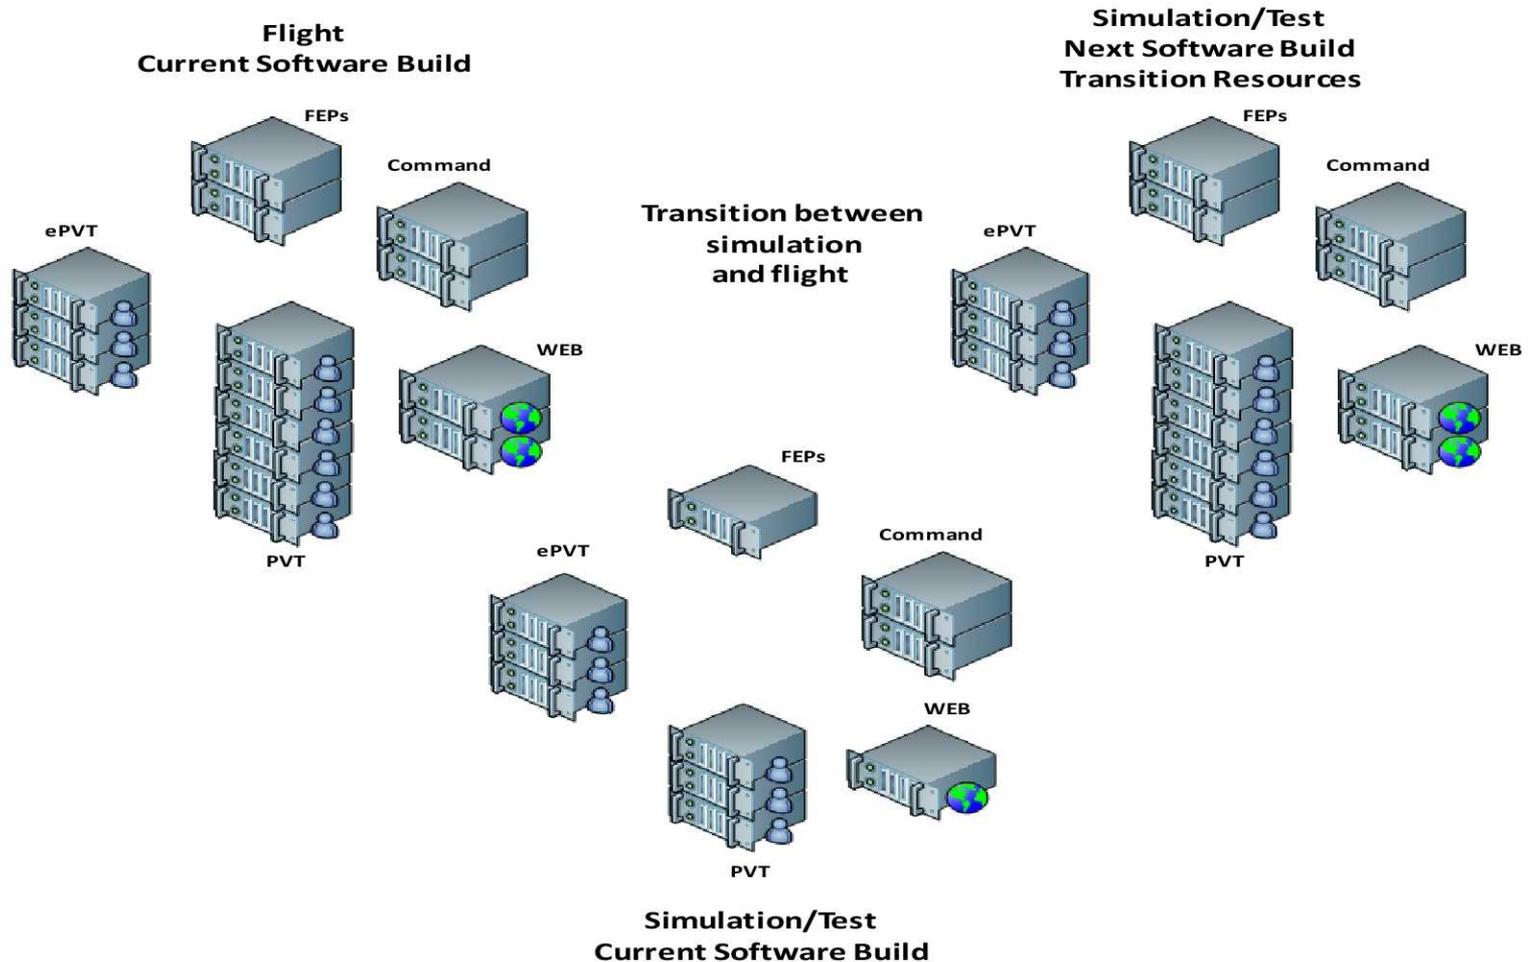


- Current HOSC Server Allocation
- Operational Considerations
- Virtual Platform Configurations
- System Configuration
- Defining a New Architecture

DEPLOYING VIRTUAL SERVERS

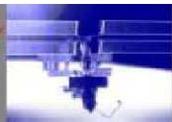


Current HOSC Server Allocation



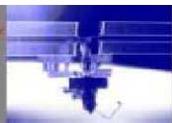
Operational Considerations

| General Consideration | HOSC Implementation |
|--|---|
| High availability power | Redundant power supplies sourced to redundant power sources (circuits) which can individually supply a server. |
| Onboard disk storage and its survivability and separability (RAID usage) | Total disk capacity is 4/73 GB drives - 10GB system disk, 24 GB allocated to each virtual server for user storage. Single Redundant Array of Disks (RAID) 10 disk configuration. |
| Physical Separation | 30+ meters to limit failures related to environmental coupling, e.g., air flow, fire suppression shower heads, loss of a circuit panel. |
| Onboard RAM and organization | Dual quad core CPU's, 24 GB total host memory |
| Network accessibility and availability of critical networks to real-time | Details on page 13. |
| Virtualization characteristics "hardened" | |
| Method of Utilization - Mission boundaries affected by virtualized platforms | HOSC Operational Concept, no total support on one platform. |
| Physical Server Configuration | Details on page 14 Red Hat™ 5.x LINUX operating system and various COTS products and locally developed custom software. |



Operational Considerations

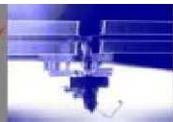
- Operations concerns
 - Platform status, both logical and physical
 - Isolation of logical platform performance characteristics
 - Ability of on-shift personnel to manage and respond to anomalies without engineering support
 - Right platform for the right mission; do not waste or restrict resources



Virtual Platform Configurations

- New virtual platforms were configured in a rough approximation of the current operational systems

| Component | Current Operational | Physical Host | Virtualized Host |
|-----------------|---|----------------------------|--------------------|
| Motherboard | SM X5DL8 | SM X7DB3 | N/A |
| Processor | 2x Intel® Xeon® CPU 3.06GHz Cache: 512 KB | 2x Intel® Xeon® X5355 CPU | 2x |
| Processor speed | | 2.66 GHz Cache: 4 MB | N/A |
| Platform cores | 1x2 | 2x4 (2) | 2 |
| FSB | 533 MHz | 1333 MHz | N/A |
| Memory | 4 GB | 24 (4) GB | 4 GB |
| RAID | RAID 10 (37) | RAID 10 (10) | RAID 10 (33) |
| Mission network | single, 100 Mbps | 1000 Mbps channel bond (1) | Virtualized |
| Archive network | single, 100 Mbps | single, 1 Gbps | Virtualized |
| Other network | N/A | single, 1000/100 Mbps | Virtualized |
| Operating Sys | LINUX RH 4.6 | LINUX RH 5.4 (HOST) | LINUX RH 4.6 Guest |

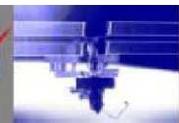
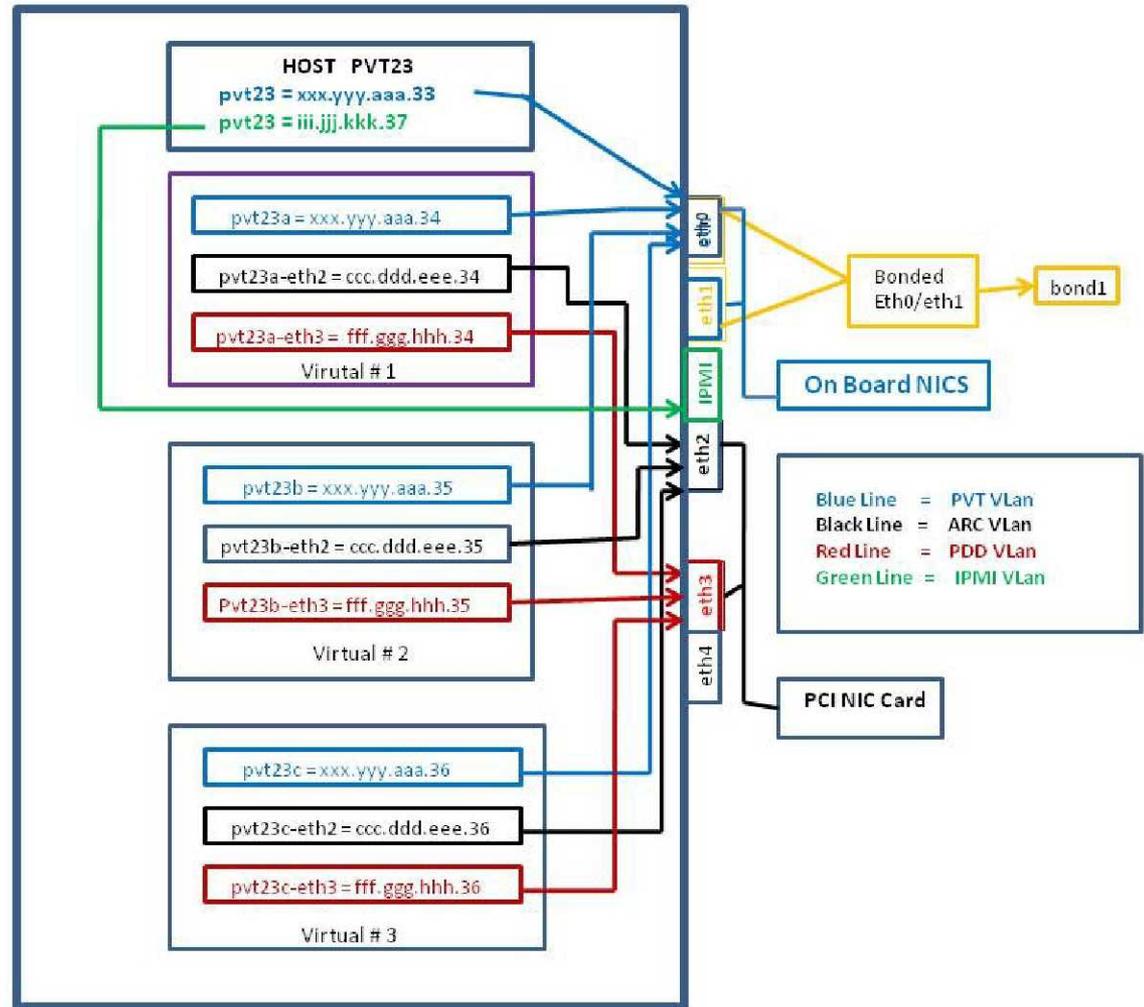




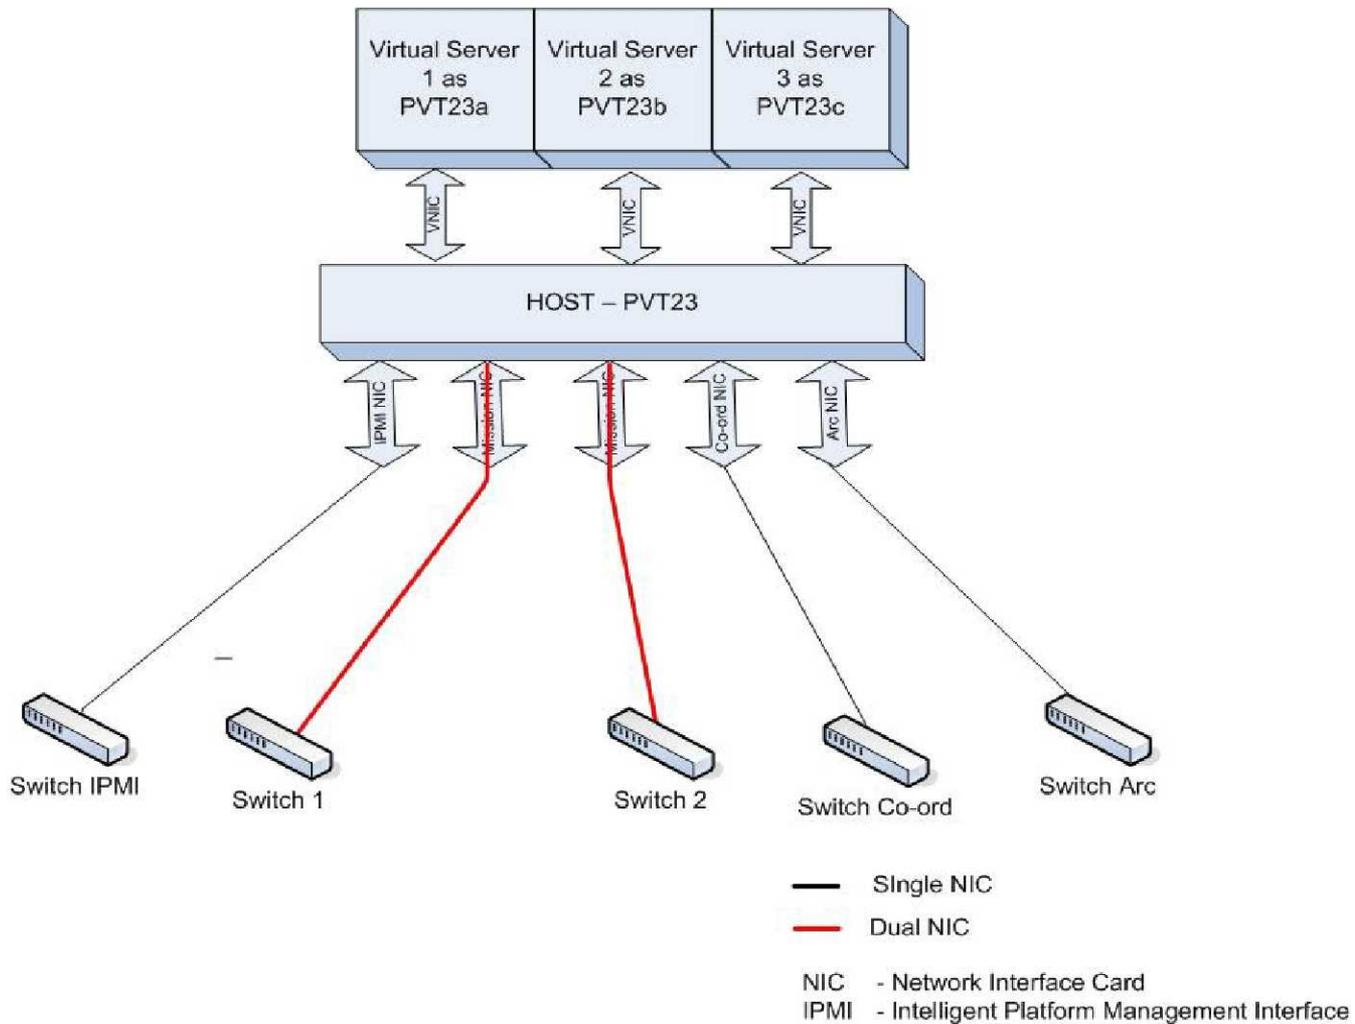
Virtual Platform Configurations

A hosted virtual = one physical server in operations with identical connectivity on Gigabit interfaces that are shared with all other virtual servers on a single physical platform.

| Subnet | Purpose |
|---|--|
| PVT VLAN | Interface that supports mission operations Composed of two channel bonded model NICs across separate Cisco switches |
| ARCvLAN | Provides access archive storage |
| Payload Distribution Device VLAN | Payload Data Distribution Network |
| Intelligent Platform Management Interface | Provides remote system management capabilities |



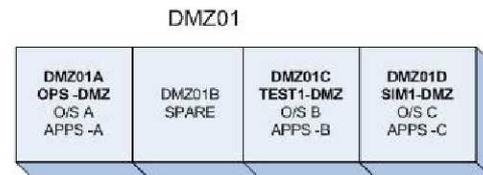
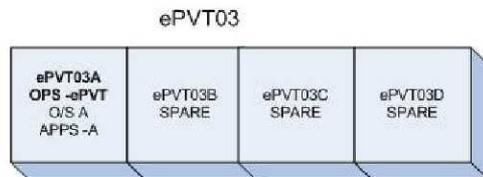
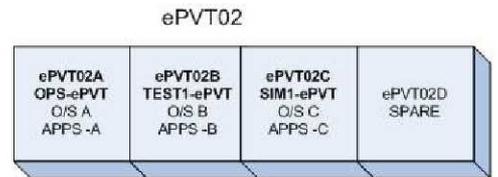
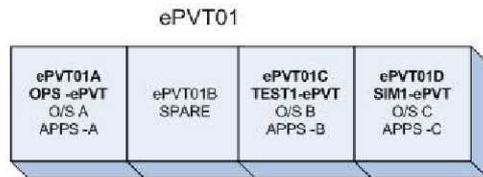
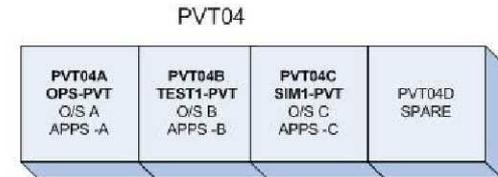
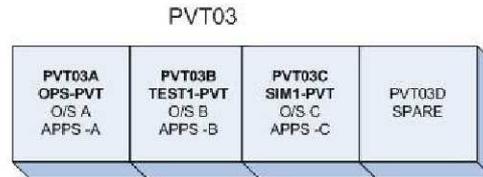
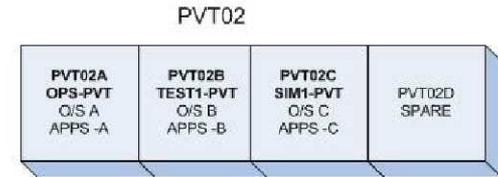
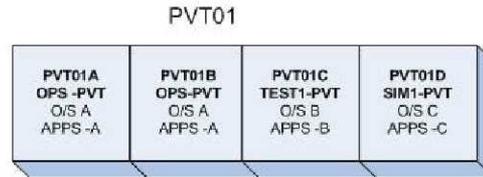
System Configuration





Defining the New Architecture

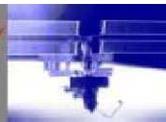
- 11 of 15 Servers are Candidates for Virtualization
- Goal – No single server outage will result in total loss of service for operations or test





- Summary of Results
- Lessons Learned

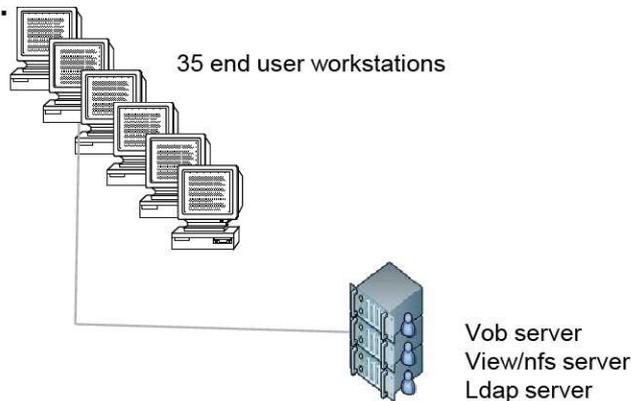
CONCLUSION



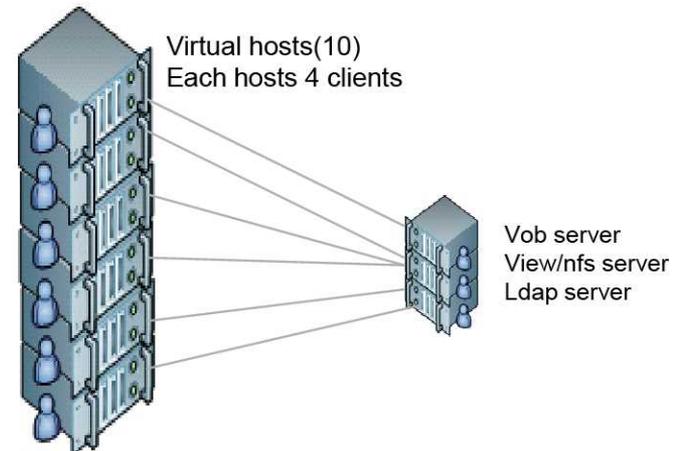
Summary of Results

- Reduced Development Servers from 35 to 10 virtual hosts; with additional system capacity. Also reduced Redhat™ software license requirements.

Prior to virtualization

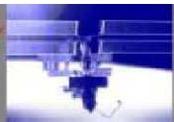
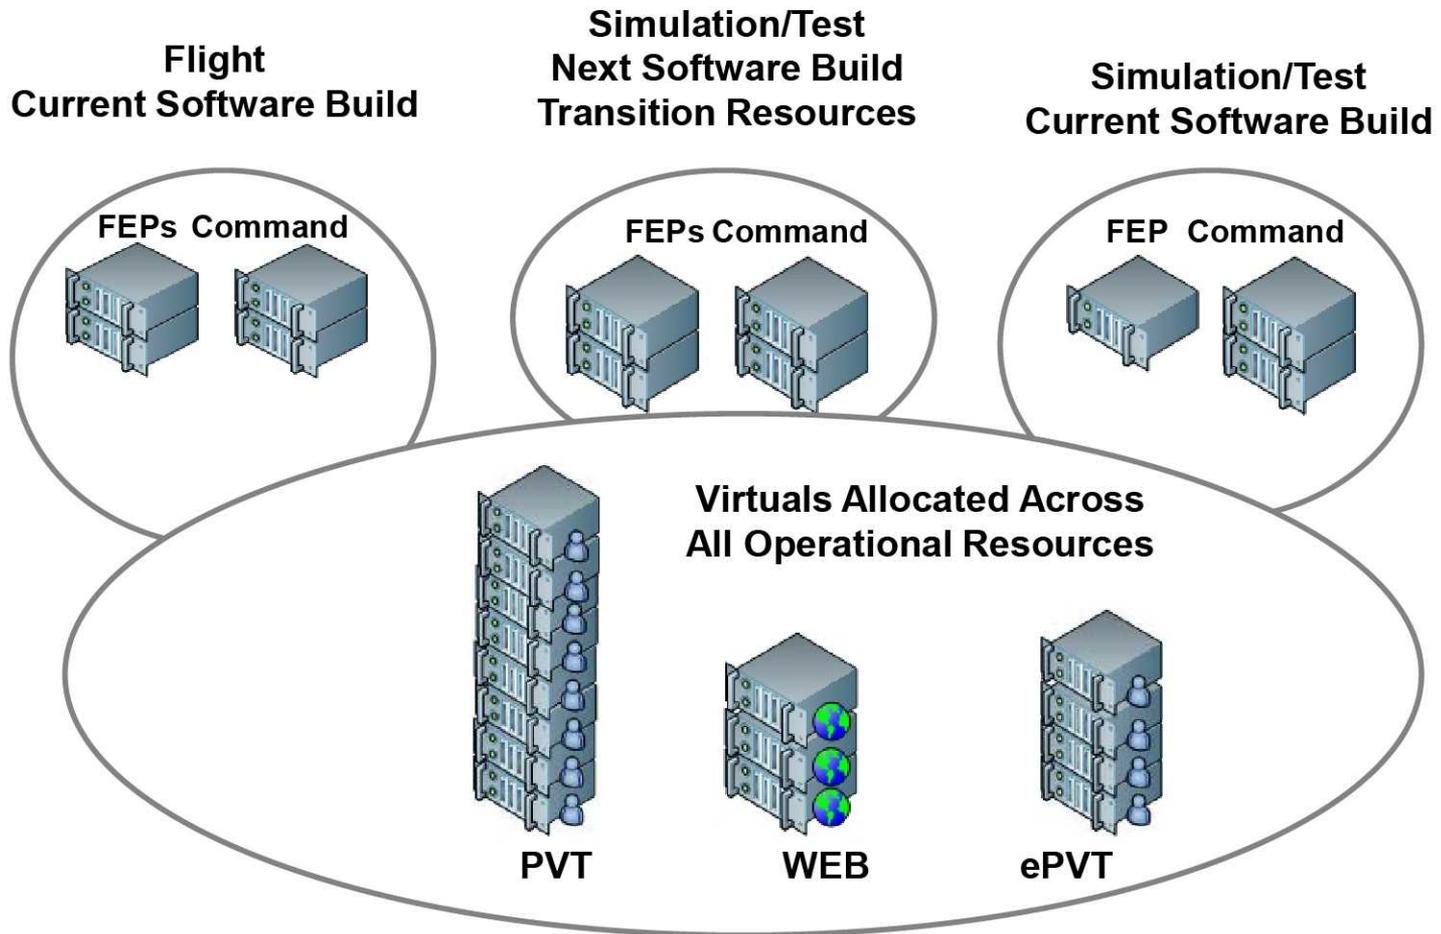


After virtualization



Summary of Results

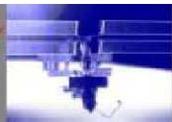
- Will realize a 59% hardware savings when all test, simulation and operations virtual servers are deployed in late 2010.



Lessons Learned



- Investigate, analyze and test the six different NIC bonding modes for the proper match to your specific use paradigm.
- New systems with max memory, drives and CPU's have larger power requirements with reduced size. Ensure you populate racks with the proper power margins, redundant circuits and adequate cooling.
- Many COTS vendors charge per CPU or core. Properly size the host CPU for the number of virtual guests you will be able to effectively use with consideration to memory, bandwidth and loading.
- Create and maintain a documented template for the host configurations and iterations you investigate so that you can easily reproduce a set of hosts.



Lessons Learned Continued

- Benchmark old and new systems so you know limitations and constraints of the virtual environment. Expect the hypervisor to use up to 10% - 20% of the CPU and memory. Size accordingly.
- Investigate the unique features that virtual machines offer before beginning your build. For example: Live migration- the ability to move a user on the fly to another virtual without interruption requires shared drives.
- Set up a test bed and test, test, test...for mission operations fault tolerance testing, use hardware failures of switches, power supplies, drives, NICs and cables. Software induced failures can be misleading.

