DISA

A Combat Support Agency

Computing Services: Virtualization

Service Design & Transition

29 July 2010

DISA
A Combat Support Agency

Topics

- State of Virtualization
- Virtualization Primer
  - CPU Usage
  - Memory Usage
  - Disk Usage
  - Customer Success Story
- Improvement to virtual architecture
  - VSphere Migration
  - Enhanced VMotion
  - Dynamic Resource Scheduler
  - High Availability
- Future Direction
  - Site Recovery Manager
  - Elastic Computing

DISA
A Combat Support Agency

State of Virtualization

- **Computing Services' Virtualization Support:**
  - **NIPR & SIPR**
  - **All operating environments, including zLinux**
  - **11 different geographic sites**
  - **Capacity Services – 3 year virtual migration**
    - **Windows/Linux**
      - vSphere hosts: 174 (NIPR), 43 (SIPR)
      - Virtual Operating Environments (VOEs): 1185 (NIPR), 246 (SIPR)
      - Allocated Storage: 34TB (NIPR), ~7TB (SIPR)
- **Next evolution is underway**
  - **VMware vSphere migration**
  - **Automation**
  - **VOE replication**

DISA
A Combat Support Agency

CPU Usage in a Virtual Environment

- **The CPUs in an ESX host are like the lanes on the streets of a European City**
  - **Large enough to handle the primary workload with ease**
  - **Frenetic, but get the job done**
- **Host is as detached from the guests as the street is from the vehicles**
  - **Streets have no idea that a car only has one passenger**
  - **Vehicles are autonomous items that run on the street**

DISA
A Combat Support Agency

Primary Workload

- Most people are just getting from one place to another
- Need no more capacity than one or two people and some bags

- 91% of CSD servers would fit into being a 1 vCPU VOE
- 66% of the VOEs are 1 vCPU, the % will grow

Our VOEs are NOT Yugos, they are Smart Cars

DISA
A Combat Support Agency

Occasional

- There are situations where things need to be larger
- Should be based on actual need and not anticipated need
- Disrupts traffic, but is necessary at times
- If you're just driving to work, don't use this

- Less than 1/3 of the VOEs are 2 vCPU
- Most were provisioned and we have not right sized them
- Should drop to less than 1/4

DISA
A Combat Support Agency

Very Rare

- Very specific uses
- Requires thought and planning

- There are only 4 production VOEs out of ~1000 that have proven a need for this
- They are Oracle databases, but we have 1 and 2 vCPU Oracle DB VOEs
- Kept as VOEs because of frequent, but short lived need

DISA
A Combat Support Agency

CPU Utilization

- Unlike the real world, we can “magically” turn a SmartCar into an RV with some evaluation and a reboot
- Start small - we can grow with a reboot
- Temporary surge capacity is easy
- Unusually high utilization will be evaluated for changes

DISA
A Combat Support Agency

Memory Usage

- Host = Bank
- Memory Size = Credit limit
- Memory Allocation = Amount drawn
- Ballooning = Requesting payments
- Swapping = Borrowing from another bank
- We're the nicest bank in the world
- Optimal state:
  - All customer's granted appropriate credit
  - We rarely ask them to make payments
  - We never borrow money from other banks
- Of CSD's physical servers, memory utilization is:

2GB or less	63%
4GB or less	24%
8GB or less	8%
>8GB	5%

DISA

A Combat Support Agency

Disk Usage

- Portable On Demand Storage (PODS)
- Box shows up on your driveway
- You load it

- No idea where it sits, other than you know it is:
  - Secured so no one else can access it
  - Protected from the elements
  - Shows up in your driveway when you ask for it

DISA
A Combat Support Agency

Customer Success Story

- 96 application servers
- Before tech refresh
  - Rates Cost (in FY10): $ 261,989
  - Software Value: $1,132,500
- If they were not virtualized
  - Rates Cost (in FY10): $ 261,989
  - Software Value: $2,265,000
- Virtualized
  - Rates Cost (in FY10): $ 173,264
  - Software Value: $ 510,000
- Savings
  - Recurring: $ 88,725
  - Capital avoidance: $1,132,500
- All are now 1 vCPU and 2GB of RAM
- Upgraded to 2003 from 2000 without ordering new hardware

DISA
A Combat Support Agency

vSphere Migration

- **VOE impact:**
  - **Must be shutdown at some point**
  - **Must be analyzed to determine right sizing, with changes to be made when the VOE is down**
  - **Must be configured to standards**
  - **Must have DRS and HA rules identified**
- **Changes:**
  - **Virtual hardware “upgrade” for VOEs (CPU, Network, etc.)**
  - **Enhanced vMotion Compatibility (EVC)**
  - **Dynamic Resource Scheduling (DRS)**
  - **High Availability (HA)**
  - **“Hot Add” of memory, on most operating systems**

DISA
A Combat Support Agency

Enhanced VMotion Compatibility (EVC)

- Vital to support for Dynamic Resource Scheduler (next chart explains this)
- Ability to mask relatively unused features in CPUs
- All hosts appear to be the same and can run any workload
- Requires the cluster to be down to perform this change
- Enabled in all new clusters

DISA
A Combat Support Agency

Dynamic Resource Scheduler (DRS)

- Balances load across hosts within a cluster
- VOEs are still bound by their same constraints
- Allows hosts to attempt to prevent ballooning and/or CPU bottlenecks
- Needs rule sets identified to maximize functionality
- Sets us up for Dynamic Power Management
- This is an enabler to allow for higher average VOE densities

DISA
A Combat Support Agency

High Availability (HA)

- Not magic
- Watches for loss of heartbeat from a virtual host in the cluster
- Executes rules that determine what to restart, when and where
- The VOE will go down. This is not 100% availability!
- Can also watch for VOE failures
- Requires network changes that will be implemented with the vSphere upgrade

DISA
A Combat Support Agency

Site Recovery Manager (SRM)

- Assists in the coordination of recovery of VOEs in a COOP event
- Policy changes must occur – single failover site
- Is not required for VOE replication
- Is very expensive
- Will require significant planning and architectural changes

DISA
A Combat Support Agency

Elastic Computing

- **Capacity (CPU, Memory, Storage) on hand to:**
  - **Absorb a reasonable number of new VOEs within the enterprise at any given time**
  - **Expand VOEs, when required, with a reboot of the OS**
  - **Sustain the loss of a host**
- **Surge**
  - **Hosts and clusters have been standardized to allow for generic augmentation**
- **Speed Of Development**
  - **Industry standard virtualization usage allows for developers to use “like” environments**
  - **Standard method in place to support a smoother transition from the developer to the DECC**


DEFENSE INFORMATION SYSTEMS AGENCY