Los Alamos “Unlimited Release” Number: LA-UR-08-2436

Classified Medialess Computing in the Applied Physics Division

Ahmad Rajeh Douglas

douglas@lanl.gov

Giving Credit Where It's Due

- **The cause of improving classified computing security was championed and funded by Robert Webster, a member of our management team, as well as John Hopson, the Los Alamos ASCI Program Director.**
- **This work was largely carried out by a four-person team:**
  - Tadeusz Raven, David Sayre, and Ahmad Douglas – Applied Physics Division
  - Scott Miller – Cyber Futures Laboratory
- **Our colleagues in the Cyber Futures Laboratory have contributed a great deal of additional work and refinement to our original solution:**
  - Mike Fisk, Alex Kent, John Parrack, Lynn Saxton, and Bill Weiss
- **The CTN X-Division Computing Support Team deployed many of the servers and services in our next-generation computing environment**
- **Anthony Clark and Danny Quist performed an in-depth evaluation of the Sun Ray hardware and software architecture**

Introduction

About the Applied Physics Division

- **Our Mission:**

  "… to apply theoretical and computational physics to the design, performance, and safety of nuclear weapons on behalf of our Nation."

- **Four hundred people, including scientists, managers, and support staff**

- **Broad Scientific Capabilities**
  - Nuclear weapon design
  - Stockpile stewardship
  - Computational analysis, simulation, and visualization
  - Nuclear threat reduction and assessment

Overview of Our IT Operations

- **User Base**
  - Scientists (require advanced visualization and computation functionality)
  - Managers (require efficient and seamless office and scheduling functionality)
  - Support (system administrators, security officers, secretaries, project leaders)
- **Back End Environment / Scale: Physically Mirrored (Open / Secure)**
  - Dozens of servers; dozens of TB of near-line storage backed up nightly
  - Vast integration issues: Solaris, Linux, IRIX, Mac OS X, and Windows
  - Local computational capabilities: Solaris, Linux, IRIX compute servers
- **The User's Office**
  - Between two and four desktop computers
  - Open Network: standard desktops
  - Secure Network: diskless (Sun, Linux, Mac, Win)
  - Average of one networked HP 4600 printer (full size color laser) per two offices

Towards A New Classified Computing Paradigm

Before We Begin: What's the Point?

- **The current model of classified desktop computing is not sufficient**
  - The end user's office remains too large of a risk
    - A "real" classified network port
    - A full-featured diskless desktop unit (PC, Mac, or Sun) with various mitigations
      - JB Weld
      - Software disablement of non-critical ports and services
      - Physically disabled Bluetooth, microphone ports, etc.
- **We need to do (much) better from the security standpoint**
  - Expand the old threat model to be more realistic
    - Yesterday: "It's impossible to stop a dedicated insider!"
    - Today: "Let's raise the bar as high as we can."
  - Manage the residual risks of desktop classified computing
    - Truly scalable security: fewer door-to-door "fire drill" mitigation exercises

Expanding the Threat Model

- **Today: The DOE cyber threat model focuses on**
  - Defending against the external attacker
    - IDS, IPS, firewalls, encryption, …
  - Curtailing the careless or unintentional insider
    - Administrative policy, media incompatibility, …

- **Can we take it to the next level? What about…**
  - Scary: the malicious insider
    - The Q-cleared Laboratory worker as our adversary
  - Scarier: the malicious system administrator
    - Q-cleared plus physical access and privileged accounts

What are the capabilities of our “new” adversaries?

- **Malicious Insider**
  - Connect hostile devices to their physically-protected office classified network port
  - Attempt to circumvent hardware and software controls on their desktop computer
  - Leverage trust relationships that may be granted to their desktop computer
  - Utilize their regular user accounts to gain unauthorized access to protected data

- **Malicious System Administrator**
  - Physical access → direct physical and cyber attacks on the server room
  - Administrative (privileged) accounts can bypass or circumvent protections
  - Able to delete or modify system logs
  - Behavior that would be suspicious from a user is within expected parameters

How can we cope with this expanded threat?

- **First Stage: Tightly control the risk exposure**
  - Consider the end-user's office to be completely hostile
  - Consider portions of the network to be completely hostile
  - Move as much risk as possible into one location: the server room
  - Enforce a "glove box" classified work environment

- **Second Stage: Harden the primary point of failure**
  - Place the server room inside of a Vault-Type Room (VTR) or Vault
  - Lock all classified computing equipment inside heavy duty custom-fab racks
  - Physically require two-person supervision for any access to the racks
    - Requires collusion of two Q-cleared insiders (much less likely)

- **Shameless Plug: Attend the Super VTR presentation!**
  - The concepts tightly integrate with our work

But, it's not just about security… can we also improve

- **… management of desktop computers?**
  - Diskless or "diskfull": many of the same problems
    - Managing the per-host configuration information
    - We're still visiting user offices for things we could manage centrally
  - Can we get to "fix it once, fix it for all (or even many / most)"?
- **… the user experience?**
  - Does a user need to have a computer for each platform?
  - Can we make using the Secure network less cumbersome?
  - "Carrots" vs. "Sticks": gaining user buy-in by improving robustness
- **TCO: the fuzzy number that keeps popping up**
  - Initial cost of hardware / software, sure… but also:
  - Maintenance, power and cooling, management cost, replacement lifecycle, space footprint, backup costs, …

Requirements Summary: Security

- **Management of edge device support and compatibility**
  - Scalable: We don't want to go door-to-door plugging ports ever again
  - Robust: Our solution has to be hardened against misconfiguration and attack
  - Adaptable: A granular solution will allow us to grant exceptions as needed

- **Raise the bar on data theft from the end-user's office**
  - The "real" classified network port
  - The network printer in the user's office (or just down the hall)
  - Attacks on the office desktop unit
    - Booting the system to a Rescue CD
    - Attaching a disk to the diskless computer
    - Connecting a microphone, video camera, etc.
    - Bridging two physically-separate LANs with multiple NICs
    - … or, think another attack up yourself! There are too many.

Requirements Summary: Usability

- **Multi-Platform**
  - Accommodate users of all major platforms: Linux, Mac OS, Solaris, Windows
- **Support ubiquitous desktop technologies on all major platforms**
  - Acrobat, Flash, Java, Microsoft Office, Streaming Media, MPEG-3, …
- **High Performance**
  - Average classified desktop machine: $5-10K (high-end visualization workstations)
  - Continue to deliver the computational and visualization power our users expect
- **Streamline the special cases**
  - One general solution should scale to meet the needs of every type of user
- **More enlightened system administration model**
  - Absolute minimum of per-client configuration state data
  - Centralize as many of the routine support tasks as possible

Researching Our Options

Tonight's Main Event: The Contenders

- **Traditional "diskfull" desktop**
  - Residual risk too high – removed from consideration altogether
  - The DOE Complex doesn't even allow these anymore!
- **Traditional "diskless" desktop**
  - Residual risk still too high – removed from consideration altogether
- **Analog KVM Display**
- **Digital KVM Display**
- **Proprietary Terminal Server**
- **Custom Terminal Server**
- **Proprietary "Virtual Display Client"**

Analog KVM Display: The Worst-Case Baseline

- **Advantages**
  - Security Mitigations: Excellent
    - Network: Analog traffic over Fiber
    - Peripherals can be restricted (HID only)
    - Stateless local unit: hard to attack
  - Platform-independent
  - Full performance of a standard desktop
  - Full dual DL-DVI graphics capability

thinklogical™
The Future of Access and Control

- **Disadvantages**
  - Large space footprint for both transmitter and receiver
  - Three dedicated fiber pairs required per unit → serious infrastructure costs
  - 1000m distance limitation seriously impairs a large classified desktop operation
  - No desktop management gains; still maintaining every desktop in the server room
  - Cost: ~$10K per TX/RX pair, plus the same high-end desktop hardware lifecycle

Digital KVM Display: Won't meet our needs

- **Advantages**
  - Security Mitigations: Passable
    - Network: IP-based KVM signals
    - Unable to find documentation on port lockdown
    - Stateless local unit: hard to attack
  - Full performance of a standard desktop
  - No distance limitation (vs. Analog KVM)

Model I9420 (Back)

Model I9420 (Front)

- **Disadvantages**
  - Not platform-independent: x86-based OSes only
  - One-to-one relationship between end-user units and server blades → we're still in the business of managing individual desktops
  - High-end units support 2x DL-DVI, but with a maximum resolution of 1600 x 1200 @ 75 Hz. That's not even in the ballpark for our visualization users.
  - Blade technology + advanced 3D requirements = open question

Proprietary Terminal Server: Won't meet our needs

- **Advantages**
  - Security
    - No local processing or disk minimizes trust
    - Peripherals can be restricted (HID only)
    - Stateless local unit
  - Performance scales with the server hardware
  - Many-to-one relationship between desktops and servers
  - Granular scaling of desktop hardware with video needs
  - More centralized system administration at medium to large scale
- **Disadvantages**
  - Network: LTSP v.4 → insecure, unencrypted protocol is unacceptable
  - We can't secure the solution manually because it's proprietary
  - Visualization requirements dictate that we use the high-end model

    → Can be attacked the same as a diskless PC

Custom Terminal Server: A Workable Possibility

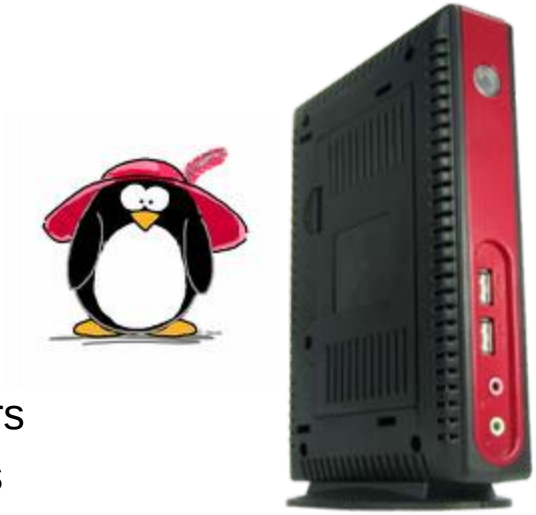

- **Advantages**
  - Security
    - No local processing or disk minimizes trust
    - Peripherals can be restricted (HID only)
    - Stateless local unit
  - Performance scales with the server hardware
  - Many-to-one relationship between desktops and servers
  - Granular scaling of desktop hardware with video needs
  - More centralized system administration at medium to large scale
- **Disadvantages**
  - LTSP integration with current RHEL is weak; extensive development required
  - LTSP v.5 alone does not provide sufficient security
  - Risk associated with a full diskless PC must be mitigated through extensive work
    - Sophisticated two-way firewall
    - Some physical attacks on the end-user terminal remain viable

Proprietary “Virtual Display Client”: More Promising

- **Advantages**
  - Security
    - Network: fully encrypted connection to server
    - Peripherals can be restricted (HID only)
    - Stateless local unit: hard to attack
  - Platform: each server runs Solaris 10 or Red Hat Enterprise
    - Supports tight multi-platform integration via “kiosk mode”
    - Many-to-one relationship between desktop units and servers
  - Servers
    - Redundancy built in with the notion of a FoG (Failover Group)
    - Hosting is implemented in software: manage one server, serve many DTUs
  - Desktop Units
    - Low cost (~$500.00), low power draw (8 W), small footprint
    - One DTU model serves basic, intermediate, and advanced users

Sun microsystems

Proprietary “Virtual Display Client”: More Promising

- **Disadvantages**
  - Tight multi-platform integration will not be trivial
    - Requires deployment of Windows / Citrix farm
    - Requires R&D of a suitable Macintosh RDP solution
  - Graphics performance still isn’t quite there
    - Dual 1900 x 1200 @ 72 Hz. is better…
    - … still not dual 2560 x 1600 @ ?? Hz.
    - What about high-end 3D?

The Ballots Are In…

- **Three possible solutions to meet our requirements**

  - Analog KVM
    - We'd rather not spend a few million dollars running new fibers
    - Even if we did, space in our server rooms would be pinched

  - Custom LTSP
    - Authoring and maintaining a custom Linux distribution still isn't our first choice
    - The necessary components are there, though
    - And our development hours would be way cheaper than implementing AKVM

  - Sun Ray
    - It's still not perfect, but can we bend it to meet our needs?
    - Even spending some time in development would be a deal
      - The Sun Ray appliance makes an attractive case vs. LTSP clones

Taking the Sun Ray 2FS to Task: In-Depth Testing

- **The data sheet sounds great… but will it deliver?**
- **"Red Team" Time: LANL has a world-class security analysis capability**
  - The Hardware: processor, flash ROM, and 64 KB of frame buffer
    - Solid security: even the flash ROM updates must be signed by Sun
  - The Software
    - Sun Ray Server Software exploit found
    - USB non-HID disablement can be attacked and defeated
      - Trivially by a privileged user! (remember: broad threat model)
- **Stress Testing**
  - Raw CPU cycles: no problem… expand the auto load-balanced FoG as needed
  - High-end Graphics
    - "Tearing" → DTU's processor can't unencrypt / decompress quickly enough
    - Native OpenGL 3D performance was very poor
    - Sun's elegant answer: Virtual GL sessions hosted on a visualization server

Operator, could you please connect us to Sun?

- **“We found this exploit in your Sun Ray Server Software…”**
  - Disclosed by Sun Microsystems and patched within two weeks

- **Can we go from 1920x1200 to full dual DL-DVI at 2560 x 1600 @ ??**
  - Sun leveraged a strategic partnership to produce an external upsampling unit
  - One SL-DVI into one DL-DVI… then multi-head groups get us all the way there
  - Verified capability within three months of request; external units ready for order
  - Functionality will be integrated into the next-generation Sun Ray DTU

- **The Sun Ray’s “tearing” precludes high-res video and advanced viz!**
  - Sun Federal requested the details of our test load and environment
  - Sun Microsystems will release a Sun Ray model with an upgraded processor
    - High performance, high resolution 3D and video @ dual 2560 x 1600
  - In the mean time, MHGs will be used to fake it to the degree possible

A path forward! The story so far…

- The Sun Ray 2FS showed the most promise to meet our requirements
- We implemented a proof of concept Sun Ray environment for testing
- Testing was positive overall, but elicited some areas for improvement
- Sun Microsystems was extremely responsive to our requests
- Sun demonstrated its commitment to our adoption of Sun Ray tech

… the Applied Physics Division then committed to a $2.0M deployment of Sun Ray technology, primarily to its production classified LAN.

Production Deployment

Stage One: Now Playing in the Server Room, it's…

- **… an all-new server infrastructure!**
  - Directory, home filespace, software licensing, software hosting, web, printing
  - Failover Groups for Solaris 10 (x86) and RHEL, SRSS kiosk servers, Citrix farm
  - Visualization farm comprised of four nVidia QuadroPlex rendering servers
  - Hardware scaling: Sun recommends 1-to-20, we did 1-to-2 (users per core)
- **… a sophisticated new network deployment!**
  - Standard Los Alamos Secure LAN routed directly into appropriate servers
  - Services LAN provides high speed server-to-server routing inside X-Division
  - Management LAN routes "Lights Out Management" (LOM) traffic
  - Desktop and Printer LAN connects appropriate servers to end-users and print rms.
  - Multi-tier service: 10GbE, 1GbE, 100MbE based on application
- **The Take Away**
  - We have been heavily entrenched in our production environment for 10+ years
  - Sometimes, you really do have to bite the bullet and start over

Stage One: Now Playing on the Desktop

- **Improved Security**
  - Truly scalable edge port lockdown
  - Solid state unit, encrypted firmware, encrypted connection: harder to attack
- **Improved User Experience**
  - Landing Zones: one desktop unit, many platforms
  - Smart Card based session mobility
  - Server-side sessions allow overnight classified processing
- **Improved System Administration**
  - Many DTUs served by a single server
  - Manage several identical servers, serve tens or hundreds of clients consistently
  - Need to scale up? Just add another server to your FoG.
    - N1 / JET system provisioning tool
  - Rich suite of system administration tools included with Sun Ray Server Software

Stage One: Supporting Roles

- **Media Assimilation**
  - Critical data from old system disks
  - Incoming media received from off-site collaborators

- **New Media Creation**
  - Data sent to off-site collaborators
  - On-site data transfer between organizations not on the Secure LAN
  - Moving between classification levels (Secret // RD up to TS // SCI machines)
  - Special applications

Stage Two: Coming Soon to a Network Near You

- **What about the Macintosh platform?**
  - You didn't miss that part… the solution is still being developed
  - Our ideal solution: AquaConnect and kiosk-mode Sun Ray servers
  - Our fallback: Analog and/or Digital KVM (with pressure to move away from Mac)

- **All end-users transitioned to the Sun Ray on the Classified LAN**
  - One Sun Ray DTU per office (more if using MHGs)
  - Screen configurations
    - 1, 2, or 4 27" displays at native resolution
    - 1 or 2 30" displays at native resolution
    - Arbitrary combinations possible – up to 32 monitors (identical displays best)
  - Improved printing experience, both in-office and at shared printers

Stage Three: About One Year Out

- **Two-way edge firewall on the User Office / Printer LAN**
  - Treats the end user offices and printer rooms as completely hostile environments
  - Prevents even a privileged user from relaxing the USB-HID device restrictions
  - Implementation: reverse engineer Sun Ray protocol, implement over ipchains
  - We can't stop every attack… but we can force most attackers to be noisy
- **Anomaly Detection**
  - Detect the noisy attacker
  - Detect users who behave outside of expected parameters
    - e.g. printing too often, working after hours, failed logins, abuse of privilege, etc.
- **Data theft rate: upper bounded by "video camera" speeds**
- **Sun N1 Grid Engine**
  - Aggregate and make use of spare FoG cycles in support of scientific work

When It's All Said and Done…

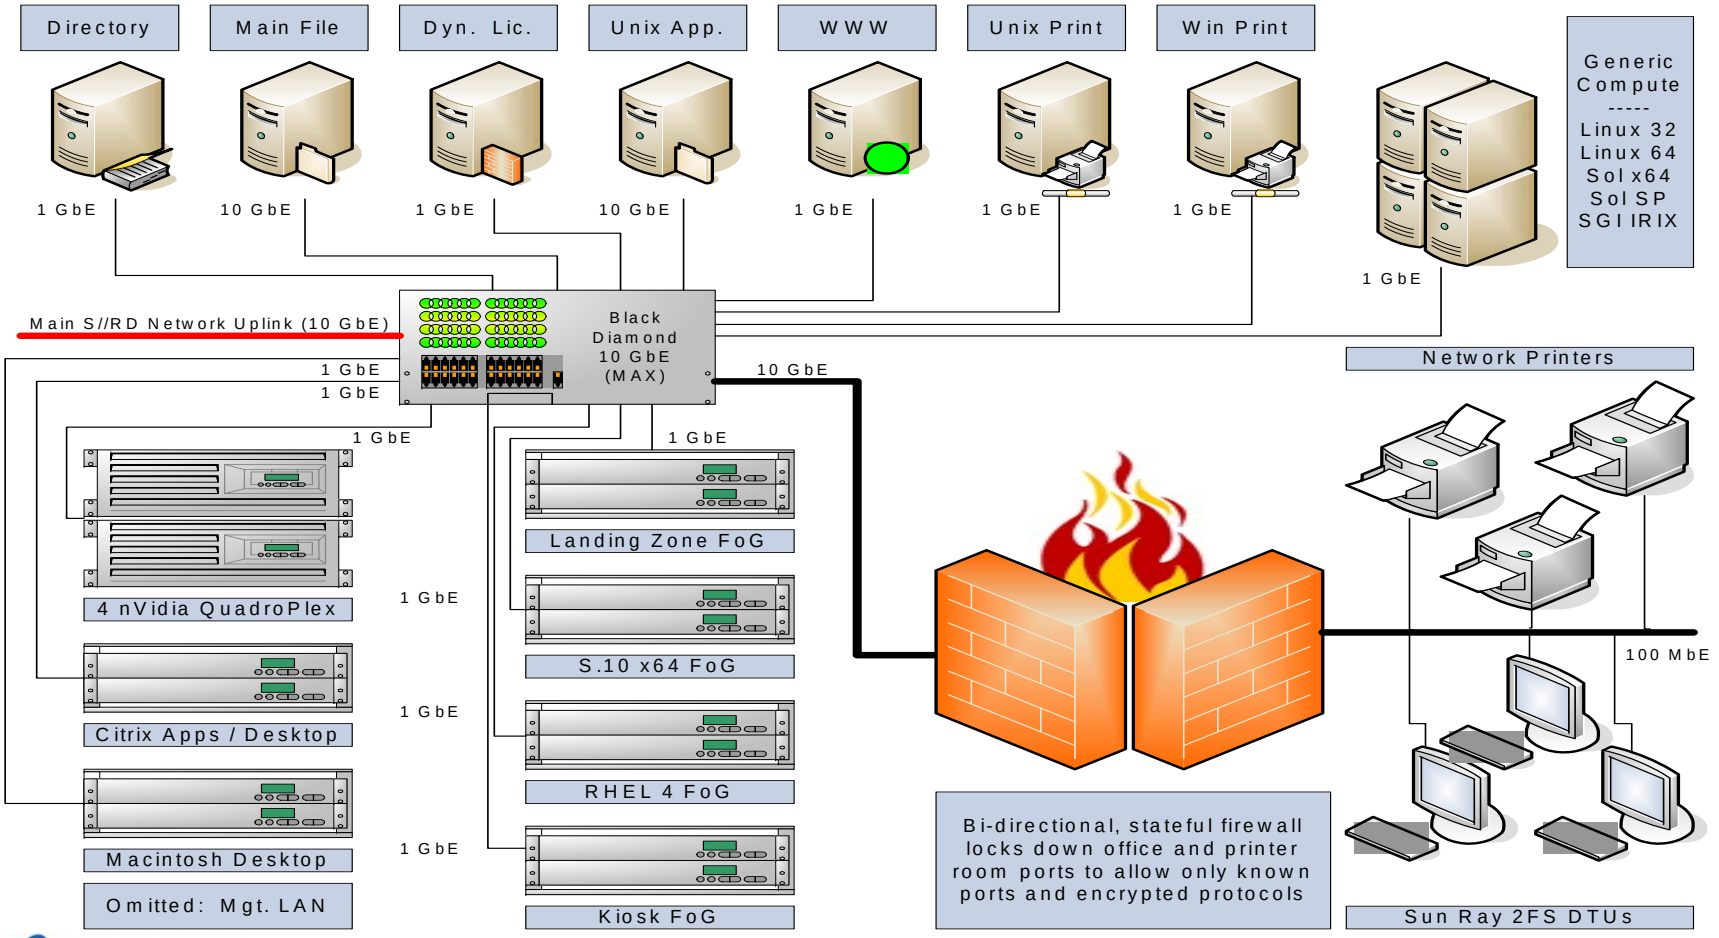

Stage Four: Into the Future

- **Put the Sun Ray platform to work on our Unclassified LAN**
  - Multiple platforms in one low-cost unit
  - Enforce media incompatibility in previously mixed-media work areas

  - Mitigate information security issues on laptops
    - VPN over Ethernet or 3G Broadband to establish a Sun Ray session
    - No Personally Identifiable Information (PII) or privileged information to lose
    - Low-cost unit reduces the impact of loss and theft

  - Improve work-at-home environment for busy scientists and managers
    - VPN over Ethernet to establish a Sun Ray session
    - No PII or privileged information to lose
    - Session mobility means your card (and session) can travel with you

  - The “Sun Global Desktop” – a software implementation of the Sun Ray client

Conclusion

Expectations

- **Improved security**
  - Immediately, from the elegant Sun Ray DTU and SRSS solution
  - In the future, from its tight integration into a hardened infrastructure
- **Improved user experience**
  - Immediately, from smart card mobility, server-side processing, centralized high-powered visualization servers, and a thorough update of our infrastructure
  - In the future, from a fresh infrastructure design, and a more scalable architecture
- **Better utilization of computational resources**
  - Central rendering servers make the best use of a fixed resource pool
  - Sun N1 Grid Engine will allow us to tackle small jobs overnight
- **Lower TCO**
  - Space / office relocation costs
  - Longer hardware replacement lifecycles
  - Fewer hardware failures (extremely simple DTU design)

Significant Energy Savings

- **[400 users] x [avg. of 2 classified computers each] = 800 desktops**
- **[800 desktops] x [avg. of 500 W each] = 400 KW power draw**
- **[400 Sun Ray DTUs] x [8 W each] = 3.2 KW power draw**
- **[100 new servers] x [avg. of 600 W each] = 60 KW power draw**
- **[400 KW old draw] – [3.2 KW + 60 KW new draw] = 336.8 KW (Delta)**
  - Assume uptime only during workday: [10 h / d] x [~250 d] = 2500 h / year
  - [336.8 KW total draw] x [2500 hours / yr] = 842 MWh per year saved…
  - At 15 cents per KWh, that's about $125,000.00 saved every year!
- **These figures are abbreviated, but conservative**
- **Also, factor in the cost of cooling eight hundred 500 W desktop computers in the height of the New Mexico summer!**

What does $2,000,000.00 buy, anyway?

- **400 high-end workstations @ $5000.00 each, *OR* …**
- **600 Sun Ray Desktop Units (DTUs)**
- **Over 100 servers to support those DTUs**
  - RHEL FoG, Solaris 10 FoG, Kiosk FoG, Windows / Citrix
  - More than 200 TB of on-line mass storage
  - Four nVidia QuadroPlex rendering servers
- **A new network to connect it all**
  - One fully-loaded Black Diamond 10GbE network switch
  - Thirty smaller supporting switches
- **… wait, there's more!**
  - Four months of on-site, full time support from a Sun Microsystems engineer
  - Four months of on-site, full time support from a Citrix-licensed engineer
  - All of the software licenses required to assemble our production environment
  - Custom fabricated high-security computer racks

The End

Questions?