

Spring 2009 [Number 243]

---



## Major Articles

NIH Federated Identity Service

The NIH Biowulf Cluster: Supercomputing for Intramural Scientists

Designing Conference Rooms for an Interactive World

WYLBUR Will Be Retiring

Introducing CIT's Updated Microsoft Office SharePoint Server (MOSS) Service

When Was the Last Time You Looked at Titan?

---

Published By  
Center for Information Technology  
National Institutes of Health  
Health and Human Services

---



# Table of Contents

---

## Articles

- NIH Federated Identity Service** 3  
*Would you like to streamline interagency cooperation? Are authentication processes hampering your embrace of Web 2.0 tools and data sharing? Find out about our award-winning Federated Identity.*
- The NIH Biowulf Cluster: Supercomputing for Intramural Scientists** 7  
*Ten years ago, a cluster of ordinary desktop computers arrived at NIH to help solve complex biomedical problems. Find out more here.*
- Designing Conference Rooms for an Interactive World** 12  
*Never again walk into your conference room faced with the hassle of cable coils, confusing buttons, or multiple remotes. Get a better conference room custom-designed for your needs.*
- WYLBUR Will Be Retiring** 16  
*WYLBUR, the Titan editing and batch processing system, will be retiring on December 31, 2009.*
- CIT Participates in IPIC 2009 Conference** 16  
*At this year's IPIC conference, the NIH Federated Identity Service received the GITEC 2009 Project Management Excellence Award and was the focus of CIT's exhibit.*
- Introducing CIT's Updated Microsoft Office SharePoint Server (MOSS) Service** 18  
*If you haven't tried MOSS yet, now is a great time: CIT is enhancing the existing service with additional features not available in the initial service offering.*
- VideoCast Website Offers a New Feature: iCalendar** 22  
*An easy way to keep track of upcoming world-class lectures and seminars - to be watched from the comfort of your own office.*
- Going Green by Increasing Print Efficiency** 27  
*Changing the way you view your Titan output can help conserve energy and save you money.*
- When Was the Last Time You Looked at Titan?** 33  
*Did you know that far from being an outdated platform, mainframes are actually proving their true value? Check out the reliability, security, and lowered costs of CIT's Titan (z/OS).*

---

**Ask the NIH Help Desk**

38

*Is your Outlook mailbox over its size limit? Create a Personal Folder to store your messages and free up space in your mailbox.*

**Directories and Reference Information**

46

**Major Contributors****Inside Back Cover**

**http://www.nih.gov** is one of the most frequently visited federal government websites.

|                                 | <i>December</i> | <i>January 09</i> | <i>February</i> |
|---------------------------------|-----------------|-------------------|-----------------|
| Total hits for the month        | 56,574,261      | 66,075,231        | 68,266,999      |
| Hits per day                    | 1,824,976       | 2,131,459         | 2,438,107       |
| Different individuals per month | 2,359,137       | 2,920,162         | 2,995,517       |

The server has been up 100% of the time\* during March 2009.

\* *Server uptime is independent of network accessibility.*

# Articles

---

## NIH Federated Identity Service

Would you like to work more closely with extramural organizations but find the exchange of information hobbled by separate security standards and multiple authentication processes?

If your willingness to share access to NIH systems, data sources, and collaborative online tools with non-NIH partners has been dampened by the hassle and cost of setting up, maintaining, and authenticating separate user accounts, consider this scenario:

A researcher from Johns Hopkins University wishes to participate in a scientific wiki created by an NIH IC. She simply uses her own login credentials from JHU to access the site – without first undergoing a lengthy credentialing process to prove her identity to NIH. There is no separate account to maintain and no new password to remember – secure shared access across institutions is easy and seamless.

If this sounds highly desirable but unrealistic, here is some good news. It is a reality now, thanks to the NIH Federated Identity Service.

### **What is the NIH Federated Identity Service?**

NIH Federated Identity Service aims to streamline interagency collaboration between NIH and outside institutions, such as universities, other Department of Health and Human Services (HHS) Operating Divisions (OPDIVs), and other federal agencies. By allowing authorized individuals to use a single name, password, or other identity to access multiple applications or data sources, collaboration is both secure and seamless.

Specifically, NIH Federated Identity Service cuts out a cumbersome separate authentication process usually required for cross-organizational cooperation. Instead of verifying user IDs separately through NIH, Federated Identity grants members secure access to NIH applications and systems by relying on the outside organization's authentication process and accepting the same user name, password, or other personal identification already authenticated by the user's home organization.

To enable this kind of shared access, the service relies on open industry standards and/or openly published specifications, leverages existing technology and infrastructure, such as Web 2.0 tools, and supports and promotes interoperability.

### **How does it work?**

NIH Federated Identity Service uses the concept of "federation" to solve the problem of separate identity authentication standards that complicates sharing access among different groups. Federation refers to an association of organizations that agree to a shared set of policies, definitions, and security standards in their authentication processes. Each member of a federated group acts as an identity credential provider

---

for its own users, but, also accepts the user authentication process of the other members when granting access to the users of the “federated” organizations. In this way, NIH Federated Identity Service maintains user privacy, by keeping the users’ credentialing process within their home organization, while enabling more seamless collaborations and transactions between federated organizations that can trust each other’s identity authentication.

How does this work for the individual users accessing applications or systems under Federated Identity? It’s easy: the Johns Hopkins-based researcher mentioned earlier simply visits the shared wiki page hosted by NIH, and, when logging in, selects her home institution from a drop-down menu. Then, since JHU is federated with NIH, she logs in using her home institution-assigned username and password. It’s that simple.

### **Who is using it?**

As of November 2008, the NIH Federated Identity Service served 16 applications at NIH. There are additional applications in the development stage for federation that are starting production in early 2009. Users at other institutions, which include universities and colleges such as Johns Hopkins University, Duke University, and Ohio State University, can interact with these applications with their own organization’s login. Potentially, about 2.2 million users from InCommon member institutions could be authorized to use NIH applications.

This year, the National Center for Research Resources (NCRR) plans to make its Clinical & Translational Science Awards (CTSA) Wiki one of the first NIH systems to be federated with non-government institutions. They will use NIH Federated Identity Service to ease access to the wiki – an online, restricted-access, collaborative work environment for members of the CTSA Consortium that currently supports 1,200 members at 38 universities. The switch is scheduled to take place in the first quarter of the year.

Among other NIH services accepting external credentials through NIH Federated Identity Service are: the new website “NIH Network for Public Information Officers (PIOs),” providing a shared workspace for communications professionals at NIH and their colleagues at NIH-funded institutions, and FDA ITAS Time and Attendance, which is hosted in the NIH Data Center.

Additional pilots of the service, readying for production this year, include the electronic Research Administration (eRA) and the NIH Portal.

### **Why use the service?**

Along with providing more seamless collaboration between NIH stakeholders and outside users, NIH Federated Identity Service also offers the following advantages:

- 
- Reduces IT expenses – NIH Federated Identity Service removes the need to set up and maintain extra user accounts for outside users;
  - Frees up Help Desk resources – Your IT Help Desk no longer has to support extramural users who have forgotten their password (again);
  - Eases access for users – Access through NIH Federated Identity Service means that users have one less password/userid to remember;
  - Lessens incidences of forgotten or expired passwords – Because the password they are using is the one they use regularly in their home organization, infrequent visitors to a federated system no longer have to worry about encountering an expired password they now can't remember;
  - Preserves user privacy – Because federation relies on the authentication process of each user's home identity credential provider, other organizations have no reason to delve into the user's personal information – it stays within the home organization;
  - Eases the use of Web 2.0 tools – NIH Federated Identity Service enables seamless but secure access to shared applications or systems, encouraging the use of Web 2.0 tools to promote secure information sharing and collaboration;
  - Offers high level of portability – NIH Federated Identity Service adheres to widely accepted, shared principles, so it is not tied to any specific protocol, technology, implementation, or vendor.

## **Measureable results and evidence of success**

The NIH Federated Identity Service has received recognition as the first federated SharePoint service in the country. In July 2008, the Federated Identity team was awarded the prestigious NIH Director's Award: "In recognition of ground-breaking work to provide improved ability to collaborate with NIH stakeholders using Federated Authentication" (NIH Director's Award Citation, July 2008). This year, the NIH Federated Identity Service was honored with the Government Information Technology Executive Council's (GITEC) 2009 Project Management Excellence Award (February 2009).

CIT receives many requests from NIH Institutes and Centers to consult on setting up the environment necessary to participate in Federated Identity. There are numerous pilot projects now underway to take advantage of the NIH Federated Identity Service for additional NIH applications. CIT also receives inquiries from other federal agencies to help set up their own Federation system, as well as requests from medical schools and research facilities to learn more about participation. Federated Identity team members are often invited to give presentations at major IT conferences, such as this year's Information Processing Interagency Conference (IPIC), where the NIH Federated Identity Service was the focus of CIT's exhibit (see also "CIT Participates in IPIC 2009 Conference" in this issue).

## **NIH service offering**

We offer the following options for the NIH Federated Identity Service:

- Participation in the service for application owners (free at NIH, cost-recovery basis for HHS and other participating federal agencies)

- 
- Consulting services (on a cost-recovery basis) for other federal agencies that want to set up their own Federated Identity Service
  - SharePoint hosting services in our secure Data Center for an outside agency's Federated Identity Service (using NIH servers on a cost-recovery basis)

### **Technical issue: data security**

The key to Federated Identity is the authentication process. The NIH Federated Identity Service depends on the ability to securely share assertions between the Service Provider, or Relying Party, and the Identity Provider. NIH has successfully used Security Assertion Markup Language (SAML) and Web Services (WS)-Federation as the protocols to facilitate the identity authentication.

**WS-Federation** provides a basic model for federation between Identity Providers and Relying Parties. FDA ITAS (Integrated Time and Attendance System), the federated project between the Food and Drug Administration (FDA) and NIH, uses WS-Federation via Active Directory Federated Services (ADFS) to share assertions.

**SAML:** NIH partners with the InCommon Federation (<http://www.incommonfederation.org/>) as a certified Identity Credential Provider. InCommon uses SAML-based authentication and authorization systems (such as Shibboleth®) to enable scalable, trusted collaborations among its community of participants. InCommon also offers identity assurance profiles that are consistent with the federal government's levels of assurance (LOA) 1 and will soon offer LOA 2, as described in NIST Special Publication 800-63. For a list of InCommon participants, see the InCommon website at <http://www.incommonfederation.org/participants/>

Both mechanisms (InCommon and WS-Federation/ADFS) broker identities and attributes, and authenticate and authorize claims, while protecting the privacy of those participating in the NIH Federated Identity Service. The goal is to provide a common infrastructure for performing Federated Identity operations for both web services and browser-based applications.

The NIH Federated Identity Service runs on Windows and Unix servers in the highly-secure NIH Data Center at NIH in Bethesda, Maryland and at an off-site disaster recovery facility.

### **Future goals**

The goal is to expand the service to use the Federal Bridge (still in the development stage) for credentialing acceptance between government agencies.

CIT plans to make the NIH Federated Identity Service a commoditized authentication service with all the necessary infrastructure and procedures defined, established, and set in place. Standardizing procedures in this way will allow quick, easy, and efficient federalization (through InCommon, the Federal Bridge, or ADFS Fed) of any additional applications.



---

## More information

You can find more information about Federated Identity at the NIH Federated Identity page (<http://federatedidentity.nih.gov>) or on the NIH Enterprise Architecture website at <http://EnterpriseArchitecture.nih.gov> or through the NIH Enterprise Architecture Community in the NIH Portal ([EnterpriseArchitecture@mail.nih.gov](mailto:EnterpriseArchitecture@mail.nih.gov)).

## Getting started

If you'd like to join the NIH Federated Identity Service, send requests for contact to [nihfederationrequest@mail.nih.gov](mailto:nihfederationrequest@mail.nih.gov).

For technical help, contact the NIH Help Desk at <http://ithelpdesk.nih.gov/support>, or by phone at 301-496-4357 (6-HELP) (local), 866-319-4357 (toll free), or 301-496-8294 (TTY).



# The NIH Biowulf Cluster: Supercomputing for Intramural Scientists

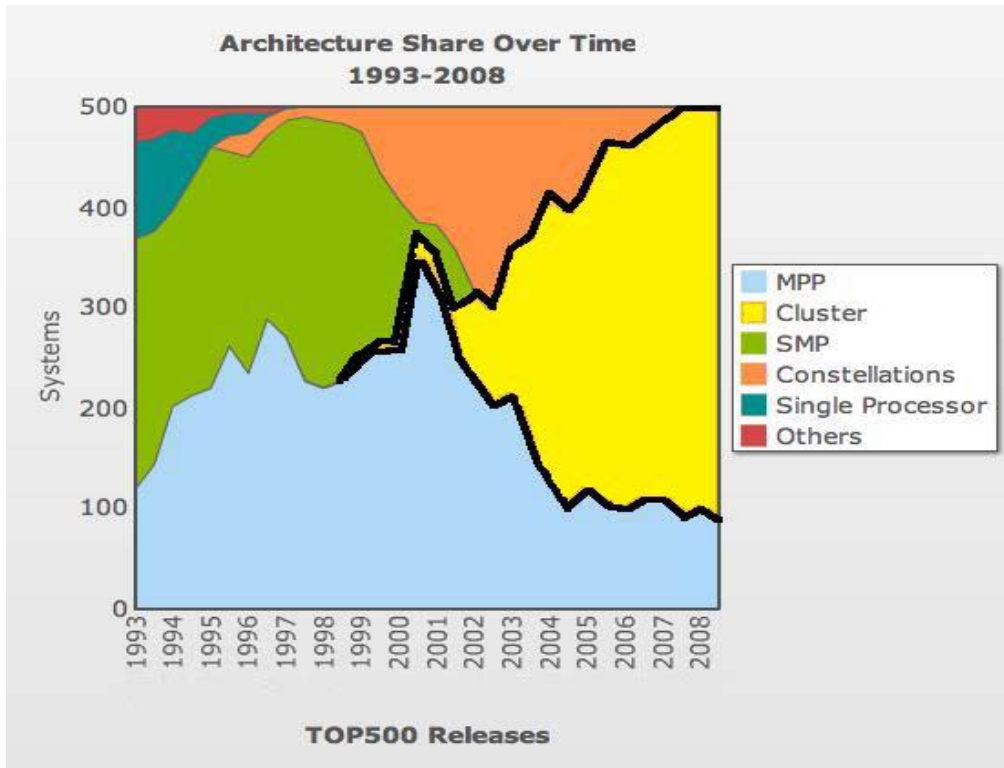
Ten years ago, a cluster of 80 ordinary desktop computers arrived at NIH for the purpose of solving complex biomedical problems. Since then, this cluster – known as Biowulf – has supported the NIH mission by providing intramural scientists with a world-class scientific supercomputing system. But where did Biowulf come from? How did it get its name? Is Biowulf right for your research project? And what is new for Biowulf in 2009?

---

## In the beginning: Beowulf

In 1994, computer scientists at NASA Goddard linked a group of inexpensive off-the-shelf personal computers together with Ethernet, creating a computing cluster with a cost-effectiveness that rivaled existing supercomputers. They called their project Beowulf.

This first Beowulf cluster was a 24-node cluster that cost \$57,000, a fraction of the cost of the commercial supercomputers then available. As off-the-shelf, or commodity, computers became even less expensive and open-source software such as the Linux operating system became available, “Beowulf-class” systems grew in popularity. In the last 15 years, as this graph from the Top 500 organization shows, clusters have come to dominate supercomputing, with 400 clusters among the world’s top 500 computers in 2008.



From [www.top500.org](http://www.top500.org), reproduced with permission

## From Beowulf to Biowulf

In 1999, DCSS’ Scientific Computing Branch went production with a Beowulf-class system of 80 nodes, which was called Biowulf – a Beowulf for Bioscience. Since then, the NIH Biowulf cluster has increased its processing power to over 6,300 processors. It is currently used by NIH intramural scientists for projects

---

ranging from molecular dynamics simulations of protein structures and genome-wide association studies, to electron microscopy image analysis and statistical calculations.



30% of the Biowulf cluster (48 processors) in 1999; and 10% of the cluster (600 processors) in 2009.

## Biowulf today

The Biowulf cluster is heterogeneous, incorporating several generations of nodes and three different interconnects. Some applications on the cluster are best suited to a particular kind of node, and because the Biowulf cluster runs many different types of biomedical applications, the heterogeneity is an advantage in the NIH environment.

Biowulf's supporting hardware includes file-system servers and a tape library. The cluster runs the Linux operating system and uses the PBS batch system for job submission and scheduling. A wide array of compilers is available for those who wish to develop programs or build code. The staff that manages and supports the cluster is comprised of both system administrators and scientists, ensuring that the Biowulf staff can address the scientific *and* the computing aspects of scientific computing.

## Parallel jobs or "swarms"?

Initially, the Biowulf staff had assumed that most jobs on the system would be parallel jobs, but within a few months, users had found a novel way to utilize this new resource: running "swarms" – large numbers of independent single-threaded jobs.

---

Biomedical projects often lend themselves to these “swarms” of computation. For example, a scientist might want to analyze 100,000 DNA sequences with a series of standard programs. Most sequence analysis programs are not parallelized, and the analysis of each sequence is independent of the next, so running them all at once as a swarm of independent jobs makes sense.

Another example is image processing: NIH scientists analyze images from electron microscopy, PET and CT scans, and MRIs, which may require independently running an image-processing program on each image. To enable easier submission of such jobs, the “swarm” program was developed in-house at NIH.

However, parallel jobs still use a large proportion of cycles on the cluster. Molecular dynamics programs such as NAMD, Charmm, Gromacs, and Amber are parallelized, and since such simulations are typically very long-running, they benefit greatly from this parallelization.

## **Publish or perish!**

The productivity of this massive computational resource is ultimately measured in scientific publications. The first publications that cited use of the Biowulf cluster appeared in 2000, about 18 months after the cluster went production. Since then, Biowulf users have continued to publish extensively, with over 80 publications in 2008 alone.



A symposium focusing on recent research citing the Biowulf cluster was held on February 3, 2009, with nine NIH researchers from diverse scientific fields speaking about the computational research they have conducted on the Biowulf cluster. Videocasts from the symposium are available at <http://biowulf.nih.gov/symposium>

## **Recent developments**

**2008: Focus on Molecular Dynamics.** Molecular Dynamics simulations of the behavior of atoms in protein structures are both computing-intensive and communications-intensive, and such jobs benefit from a high-performance network between the nodes. This type of job accounts for a significant proportion of the cpu cycles used on the cluster, so every evolution of Biowulf has included some nodes connected by a high-speed, low-latency network such as Myrinet, and more recently, Infiniband. In 2008, the entire hardware upgrade was targeted at molecular dynamics jobs, and we added almost 1800 processors connected by Infiniband to the cluster.

**2009: Focus on Storage.** In the last decade, the storage needs for many scientific projects on the Biowulf cluster have increased from 100s of Gigabytes (GB) to multiple-Terabytes (TB). For example, a study on the functional consequences of human genetic variation on brain function from NIA and NIMH involves 193 human brain samples of 13,000 transcripts in each of 4 brain regions. The initial analysis required 1 TB, and the subsequent steps will require at least 5 TB. To enable such projects, the 2009 upgrade to the

---

cluster has focused on storage, leading to an additional 500 TB of storage to be available in late spring 2009.

### **When do you need to use Biowulf?**

- If you have long jobs such as protein simulations that would take months or years on a desktop computer.
- If you have large numbers of independent jobs, such as running Repeatmasker and Blast on many thousands of sequences.
- If your jobs require large memory. For example, a Gaussian job requiring 20 GB of memory can be run on the Biowulf “fat node,” a special node which is designed for large-memory jobs.
- If your jobs can run in parallel on a distributed-memory system such as Biowulf. Many molecular dynamics programs are parallelized, and running a job on 128 processors can enable large simulations that would otherwise not be possible.

### **When should you not use Biowulf?**

- “Serial” jobs: (those where each successive process depends on the previous step) will not benefit from being run on a system like Biowulf, since such jobs cannot utilize multiple processors. Such jobs will run just as fast on a desktop system.
- Small numbers of short jobs. If your proposed project requires limited computational power, it’s not worth the overhead of running via a batch job on Biowulf. Use Helix (<http://helix.nih.gov>) instead!
- Interactive jobs. The Biowulf batch system is most useful for jobs that can run as detached processes, that do not depend on an interactive connection, and that do not require large amounts of graphic information being transmitted through the network. Thus, programs that can only run via a GUI are generally not suitable for the cluster.

### **Useful URLs**

How to get a Helix/Biowulf account: <http://helix.nih.gov/Documentation/accounts.html>

Applications installed on Biowulf: <http://biowulf.nih.gov/apps/>

The Biowulf User Guide: [http://biowulf.nih.gov/user\\_guide.html](http://biowulf.nih.gov/user_guide.html)

### **Questions? Contact us**

The Biowulf staff welcomes your questions. To ask us about the Biowulf cluster, the suitability of jobs, available resources and applications, or anything else you’d like to know about Biowulf, please send email to [staff@helix.nih.gov](mailto:staff@helix.nih.gov)




---

# Designing Conference Rooms for an Interactive World

Once upon a time, conference rooms were simply places large enough for the entire group to meet. The speaker stood at a podium and attendees strained to hear over the shuffling of paper handouts. But changes in technology have redefined meetings, and conference rooms need to evolve also. Now you might need to:

- Hold a group meeting even though people are stationed in three different buildings on two different campuses;
- Meet with a group in California, without buying plane tickets;
- Conduct a public forum without renting a hotel ballroom.



The Room Design team of CIT's Collaborative Technology Innovation and Video Services (CTIVS) can design, equip, and build a conference room that supports these or any other events your mission requires. While an initial investment is required, a modern conference room allows you to take full advantage of new technologies and will ultimately save your organization both time and money.

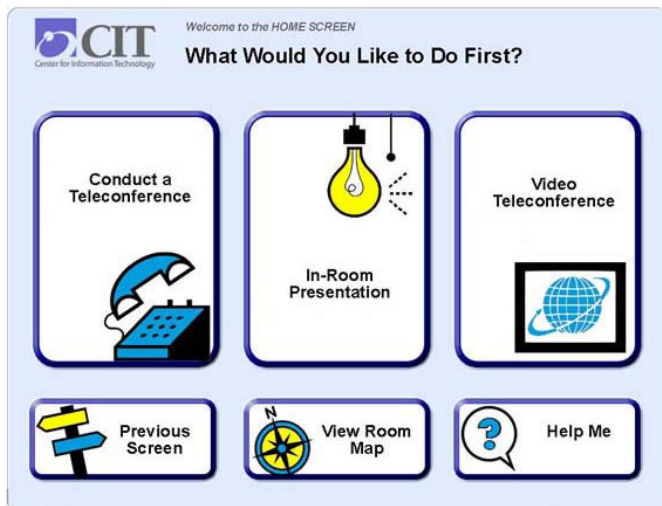
## Run a meeting at the touch of a button

Never again walk into your conference room faced with the hassle of rerouting cables, the annoyance of

remembering which buttons to press, or the frustration of finally finding the correct remote only to discover that the battery is dead.

Instead, access the room functions with a touch panel specially designed to meet your requirements.

The sample touch panel shown here allows users to choose between an audio conference, an in-room presentation, or a video teleconference, simply by touching





---

a choice. If you are holding a traditional meeting where everyone is in the room with you and needs to look at the same document, simply select “In-Room Presentation.” The system then automatically dims the room lights, closes the window shades, turns on the projector, lowers the projection screen, and displays the contents of your laptop screen onto the projection screen – all of this set in motion simply with one touch of the touch panel.

Of course your touch panel allows you to override the default settings to fit your needs. If you are video conferencing with remote attendees and want to simultaneously show a presentation slide you can use the touch panel to shift from displaying the remote attendees in full-screen mode to showing them in a small “picture-in-picture” window overlaying your PowerPoint presentation.

Meanwhile, all the equipment and cables that make your multimedia system operate are neatly tucked away in cabinets.



*[Caption: This room has twin wall-mounted plasma screens above the hardware cabinets. The touch-panel is on the counter at the far right.]*

## **Getting your own custom multimedia conference room**

Every step can be handled for you, from demolition of an existing space, to specifications, new construction, equipment installation, automation, and training.

In order to implement a system that truly meets your needs, your objectives must be clearly defined. A project leader will work with you to determine what you want the room to do. Will the room be used for in-person meetings, video teleconferencing, online meetings, or combinations of these? Will the meeting be conducted by administrators showing PowerPoint presentations or researchers discussing high-resolution medical images?

Together, we will develop specific requirements to ensure that your new conference room will meet your objectives. For example, you may want a user-friendly touch panel that enables anyone on your staff to hold a Video Teleconferencing (VTC) call, while another group prefers to lock the controls so that only authorized personnel can initiate a call. Your space may have large windows that require a brighter projector or room darkening shades, while another room needs sound barriers and special audio equipment because it is located between a high-traffic street and a noisy office break-room.

---

During the engineering stage, the project leader will determine what components will be necessary and how they will work together to meet your requirements. Typically, a multimedia room will include a projector that displays files from a laptop or a dedicated room computer onto a screen. The projector shown here is mounted on the ceiling where it is unlikely to be obscured by someone walking in front of it. This placement also discourages people from adjusting the settings.



Rooms usually include at least one digital video camera to capture video images of the people at the conference room table or the speaker at the podium. The video is sent to a group you are conferring with in a distant conference room, which is sending their own video back to your room. Thus, both groups can see each other and hold a much more natural conversation whether they are separated by miles or continents.



*[A Pan/Tilt/Zoom capable digital video camera mounted in the corner sends video of the in-room attendees to remote attendees via VideoTeleconferencing. Both groups can see each other on screen.]*

There is no limit to what can be created. If there is a way to make your meeting more productive, we can help you. Other options include:

- Multiple microphones on the table or suspended from the ceiling to pick up sound from everyone in the room and transmit it to the remote attendees;
- Ceiling speakers to allow you to hear the remote attendees;
- DVD players that automatically record your meeting to a disc ;
- "Smart boards," or electronic whiteboards, that allow you to save the information written on them as digital files.



---

Finally, with all the planning, specifications, and ordering complete, the physical installation of your new conference room can proceed. Once your room is ready for use, we can train your personnel on how to use it, and provide ongoing maintenance so your room will always be ready for the next meeting.

### **For more information**

With CTIVS acting as your consultant and general contractor, you can reap the benefits of new communication technology with minimal effort on your part.

We look forward to discussing how your group can acquire a user-friendly, multimedia conference room that allows you to share knowledge without wasting time or money traveling.

Contact us directly:

CTIVS Room Design and Support Team

(301) 594-8433

Email: [CITVideoRoomDesignSupport@mail.nih.gov](mailto:CITVideoRoomDesignSupport@mail.nih.gov)

Website: <http://video.nih.gov>

Or contact us through the NIH Help Desk at:

(301) 496-4357 (6 HELP) (local)

(866) 319-4357 (toll free)

(301) 496-8294 (TTY)

<http://ithelpdesk.nih.gov>



---

## WYLBUR Will Be Retiring

WYLBUR, the Titan editing and batch processing system, will be retiring on December 31, 2009. Many of WYLBUR's text editing functions can now be performed with desktop computer tools. In addition, we have many off-the-shelf products on Titan that will provide customers with the same functionality as WYLBUR command procedures.

We will do everything we can to help you with the transition, including working with you to convert your command procedures to REXX EXECs (an interactive language for TSO). For customers who use WYLBUR interactively, we will be offering classes in ISPF (Interactive System Productivity Facility), a full-screen editor in TSO.

We want to hear from you and make sure we address your concerns. Please contact the NIH Help Desk at <http://ithelpdesk.nih.gov/support> or 301-496-4357, 301-496-8294 (TTY), or toll free at 866-319-4357 with your questions and requests for assistance. To keep up with the latest news on the WYLBUR retirement, subscribe to the *Titan News* online newsletter at <https://list.nih.gov/cgi-bin/wa?SUBED1=cit-titan-news&A=1>.



## CIT Participates in IPIC 2009 Conference

The Center for Information Technology sent representatives to this year's Information Processing Interagency Conference (IPIC) and set up an exhibit featuring the NIH Federated Identity Service (<http://federatedidentity.nih.gov>) and other services hosted at the NIH Data Center. The conference was held in Orlando, Florida, from March 1 - 6, 2009.

IPIC (<http://www.ipicconference.com/9IP/>), sponsored by the Government Information Technology Executive Council (GITEC) (<http://www.gitec.org/>), brings government and industry IT executives together to share and collaborate on IT innovations, challenges, and successes. IPIC is a respected annual forum for government leaders, industry, and academia to share best practices regarding Information Technology.

---

## CIT's participation

CIT's NIH Data Center was one of many federal data centers and IT companies participating in the exhibit hall of this year's conference. The CIT exhibit highlighted the NIH Federated Identity Service, which was awarded a prestigious GITEC 2009 Project Management Excellence Award at the conference. Visitors to the CIT booth expressed genuine interest in this project, which allows authorized individuals to use the login from their own organization to seamlessly access NIH applications, such as SharePoint sites. (For more on this service, please see the "NIH Federated Identity Service" article in this issue.)



Valerie Wampler, co-manager of the project, was at the conference to accept the award on behalf of her co-manager, Debbie Bucci, and the other members of the CIT/NIH team. The award letter noted that the "NIH Federated Identity Service exemplifies outstanding project management excellence in the award category of Enhanced Business Results from Innovation."

Adriane Burton, the Director of the Division of Computer System Services (DCSS) within CIT, moderated a panel discussion on "Success Stories and Lessons Learned on Web 2.0 Technologies." Valerie Wampler from CIT/DCSS and Adrian Gardner, the Chief Information Officer from the National Weather Service, presented on Federated Authentication and Web 2.0 technologies used at the NIH and the National Weather Service. Some of the challenges that were identified were: governance, authoritative sources for information, and training.

In addition to Valerie and Adriane, other CIT/DCSS attendees were Norma Stern and Chantel Walker Steward.

## The conference theme

The theme for this year's conference was "Transcending Transition, Transcending Technology: A First Look at the Impact on Government Information Systems Programs and Technology as We Transition into

---

the New Presidential Agenda" and included programs on the changing budgetary and IT landscape, Web 2.0, innovation in government, and cyber security.

### **Conference activities**

The speakers and panel members included Kevin Plexico, from Input, on the impact of the government stimulus package; David Pogue, Personal Technologist, of the *New York Times*, on Web 2.0; David Evans, from Cisco, on a look at future in IT; John Gilligan, from Gilligan Group, Inc., who stressed the need to pay more attention to cyber security; and other IT executives from the private sector and government.

In addition to the government and industry exhibit booths and the excellent speaker and panel discussion programs, there were social/networking events that raised funds for the featured charity, the Brain Injury Association of America (BIAA).



## **Introducing CIT's Updated Microsoft Office SharePoint Server (MOSS) Service**

If you haven't considered using CIT's SharePoint Service offering yet, you might want to take another look. We are pleased to announce a modification to our existing SharePoint service.

When we launched CIT's collaborative service, we used Microsoft Windows SharePoint Services (WSS) to give customers flexible team sites for easier information exchange and shared productivity.

Starting in the 2nd Quarter FY2009, we are enhancing the existing service with additional features not available in the initial offering. The most notable additions include content management services, access to premium MOSS features, and an infrastructure upgrade to include an SQL clustered database solution.

---

## What is MOSS?

MOSS is a web service that provides an organized hierarchy for information sharing, document collaboration, and content management, allowing for easy and efficient cooperation on common projects. Teams can edit documents, assign tasks, publish content, plan events, and attend meetings regardless of the geographical location of individual team members.

With the release of Microsoft Office SharePoint Server (MOSS), Microsoft integrated and consolidated SharePoint and Content Management Service (CMS) into a single service. CMS is a publishing function that allows content creators, or “subject matter experts” (SME), to publish their own web page content without the intervention of the web administrators.

As a result of consolidating SharePoint and CMS, the MOSS service now provides customers with both the collaboration component from SharePoint and the content management component from CMS. If you are a current SharePoint customer, you will also have access to the CMS features.

**Please note:** With the release of MOSS, CMS 2002 will be phased out. CIT will be working with our current CMS customers on their migration to MOSS.

## Collaboration with MOSS

MOSS sites’ collaboration spaces make sharing ideas and contributing content easier for teams. The meeting workspace offers more than a place to meet online; your team can also save and share agendas and minutes, goals and objectives, attendees’ contact information, and other meeting-related documents in the workspace. If you often collaborate on documents, consider the benefits MOSS sites provide for shared document management:

- **Check-in and check-out:** Provides users with exclusive rights to a document. When a document is checked out, only the user can make changes and save the document. This gives multiple users a way to collaborate on the document without overlap in work.
- **Version control:** Helps users track changes, merge versions, and rollback versions if necessary.
- **Content Approval:** Allows authors to configure content approval so that users are required to obtain approval on content before posting documents.

## Features that make MOSS services easy to use

**Compatibility:** Microsoft Office SharePoint Server 2007 is integrated and compatible with Microsoft Office 2003 and Office 2007 for Windows. It is also compatible with Microsoft Office 2004 for Macintosh and multiple browsers.

**Convenience:** If you use Microsoft Office programs, then you are already familiar with the basic tools for accessing SharePoint features. The integration of SharePoint with Microsoft Office for Windows (see

---

“Compatibility”) means that you can use programs like Outlook and Word to take advantage of MOSS collaborative functions:

- **Meeting Workspaces:** Create meeting workspaces directly from Microsoft Outlook. Set alerts to notify users, by email, when information on SharePoint changes.
- **Document Workspaces:** Create document workspaces directly from Microsoft Word to write, edit, and share documents.
- **Task Panes:** Access Task, Link, and Contact Lists through Office 2003 and Office 2007 task panes.
- **InfoPath Workflow:** SharePoint Services are integrated with Microsoft Office InfoPath 2007, so you can use InfoPath 2007 capabilities to design and publish interactive, user-friendly form templates.

**Custom NIH templates:** In order to make setting up your SharePoint site easier, CIT has created custom NIH templates with a standard NIH header and footer. This way, you can choose which template best fits your needs. The template options are tailored for team sites, document sites, meeting sites, content sites, and discussion sites.

## Security

To customize and increase security, Site Administrators can assign levels of permissions to users that define their degree of access to a SharePoint site. Administrators can assign user groups to the following pre-defined roles:

- **Reader:** A Reader has read-only access to the site.
- **Contributor:** A Contributor can add content to the existing document libraries and lists.
- **Web Designer:** A Web Designer can create the libraries and lists, as well as customize pages in the site.
- **Administrator:** The Site Administrator has full control over the SharePoint site.

If these established roles are not sufficient, users can create custom roles as needed.

SharePoint security is hierarchical, distributed, and granular:

- **Hierarchical:** Permissions can be granted at a root site and inherited several levels down if desired.
- **Distributed:** Site Administrators can create cross-site groups so users can have access to multiple sites.
- **Granular:** Users can be given permission to specific documents, or entire site collections.

And don't worry about remembering yet another loginID and password. Because SharePoint is integrated with Active Directory and NIH Login, the same account that you use to access other NIH resources works for your SharePoint sites too. Finally, if you want to cooperate with non-NIH partners, the Site Administrator can grant external collaborators access to SharePoint sites as needed.

---

## Publishing with MOSS

SharePoint Publishing automates the online publication process, making it easy to use your browser to create new pages, edit existing content, and submit content for review. SharePoint also ensures that all content creation adheres to the established workflow before publishing approved content.

SharePoint Publishing gives you the tools to easily and quickly publish content to your sites. As an author, you will be able to publish content from your development sites to your production site(s) without requiring help from the Site Administrators or IT department. This will help to decrease the amount of support required and make it easier and faster for you to deploy content out to the NIH community.

## CIT's Windows MOSS standard offering

CIT will offer its MOSS Service for shared and dedicated server environments in the 2<sup>nd</sup> quarter of calendar year 2009. If you already have a previous SharePoint site collection, it will be migrated to the new Microsoft Office SharePoint Server 2007 environment.

CIT has built a fully redundant, resilient, and well-designed environment for this production offering. This includes F5 Load balancing and automatic failover, 24x7 site monitoring, and an SQL clustered database solution, as well as document-level backup and restore, and options for managing site growth. SharePoint basic client access licenses (CALs) are included for customers who participate in the HHS Enterprise Agreement (e.g. iSDP).

Each customer must designate a Site Administrator who will be given administrative access to their project's site. As part of the offering, we provide the Site Administrator with information and training on how to oversee the SharePoint site. CIT conducts training sessions at least once a quarter to ensure Site Administrators have a solid background in both the SharePoint product and their site responsibilities.

## MOSS premium features

With the new Moss offering, customers who have purchased premium Client Access Licenses (CALs) can use the Excel Services, Business Intelligence, and Business Process Forms features in MOSS. Each customer is responsible for purchasing their own required CALs from the HHS iSDP support or Microsoft.

Some examples of the benefits these features offer include:

- **Excel:** Business data search and Excel spreadsheet developer reuse and extensibility;
- **Business Intelligence:** Web-based business intelligence using Excel, Integrated Dashboards, and Key Performance Indicators;

- 
- **Business Process Forms:** Browser-based forms, centralized forms management, and control and “Design Once” development model.

For an additional fee, support for Rights Management Services (RMS) will be available to customers. Organizations can use RMS to help prevent sensitive information from intentionally or accidentally getting into the wrong hands, since RMS protects information through usage policies no matter where it goes. Each customer is responsible for purchasing their own required RMS client access licenses (CALs) from the HHS iSDP support or Microsoft.

Support for custom code will also be available to customers, for an additional fee. CIT staff will provide individualized, hands-on administrative assistance required at the global level in moving custom code to the production environment.

### **How to get more information**

If you are interested in MOSS services or would like to migrate your existing SharePoint or CMS site to the MOSS platform, please contact the NIH Help Desk (<http://ithelpdesk.nih.gov>) by phone at 301-496-HELP (301-496-4357), 866-319-4357 (toll free), or 301-496-8294 (TTY).



## **VideoCast Website Offers a New Feature: iCalendar**

Attend world class seminars from your office: from in-depth lectures about cutting-edge medical research to workshops on preventing identity theft, dozens of special NIH events, seminars, and lectures are available each month from the Center for Information Technology (CIT) VideoCast website.

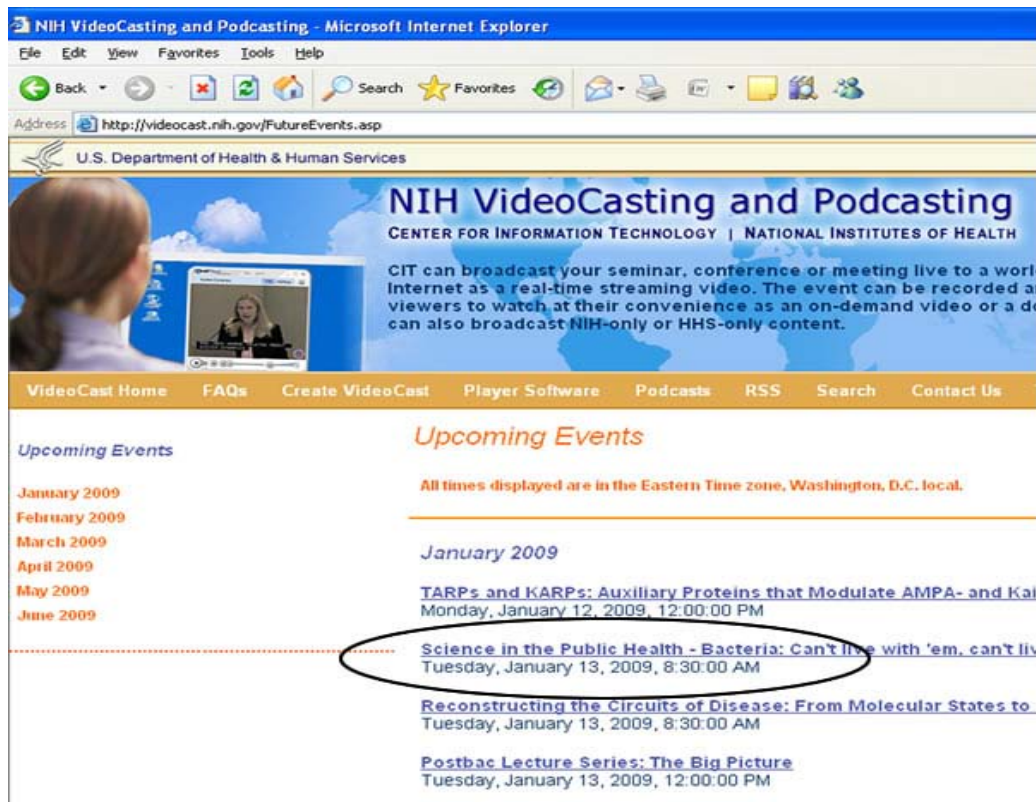
VideoCasting is a method of electronically streaming digitally encoded video and audio data from a server to a client computer. Streaming files are not downloaded, but rather are broadcast in a manner similar to television broadcasts.



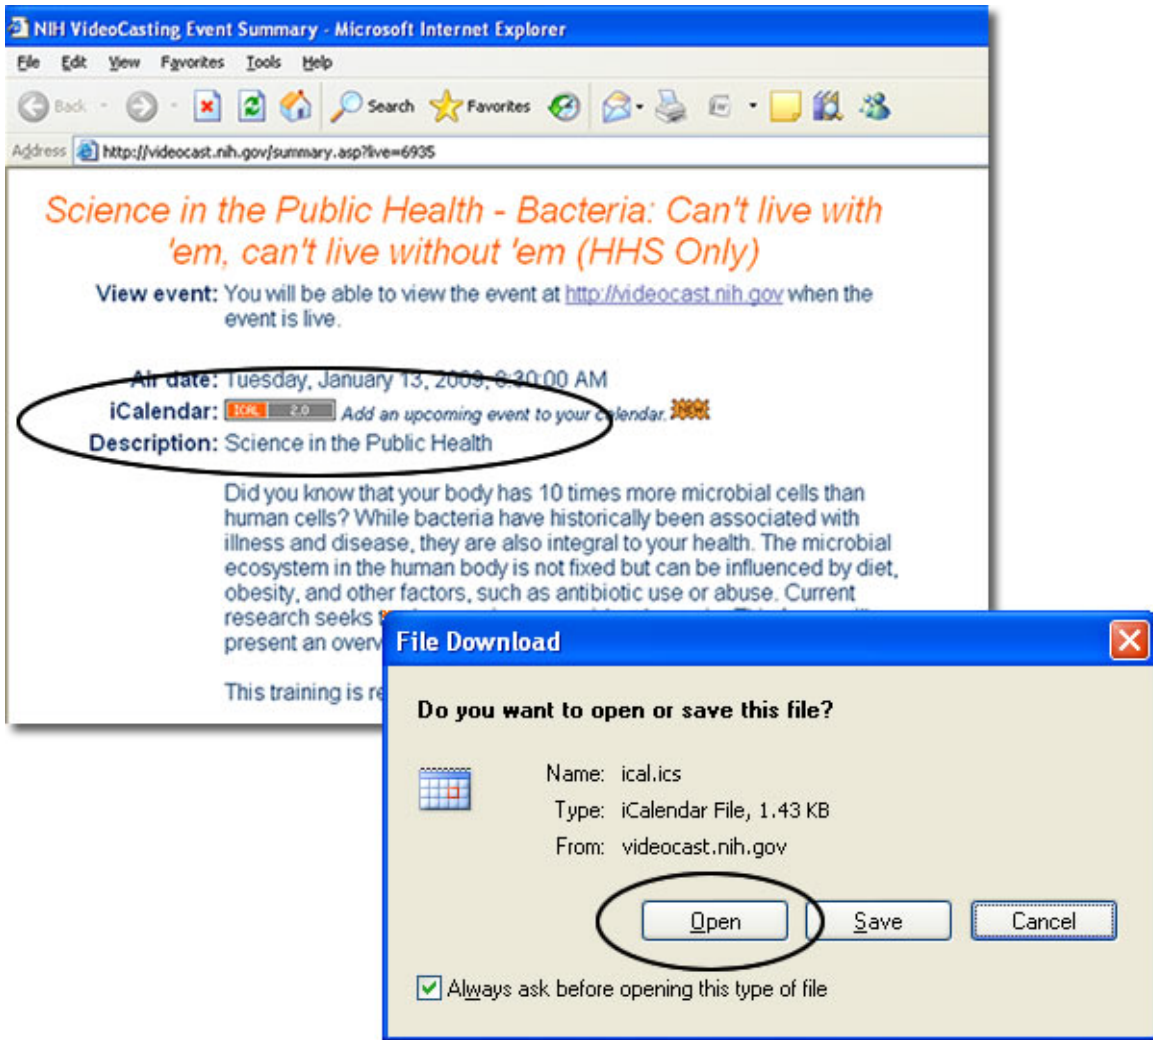
You can find a schedule of upcoming events, as well as an archive of past VideoCasts, at the CIT VideoCast website (<http://videocast.nih.gov/>). Upcoming events are broadcast online at the date and time indicated on our website. The site now also offers iCalendar, a new feature that enables you to add scheduled VideoCasts to your Outlook Calendar, so you can easily keep track of upcoming events you would like to see.

## How to use the iCalendar feature

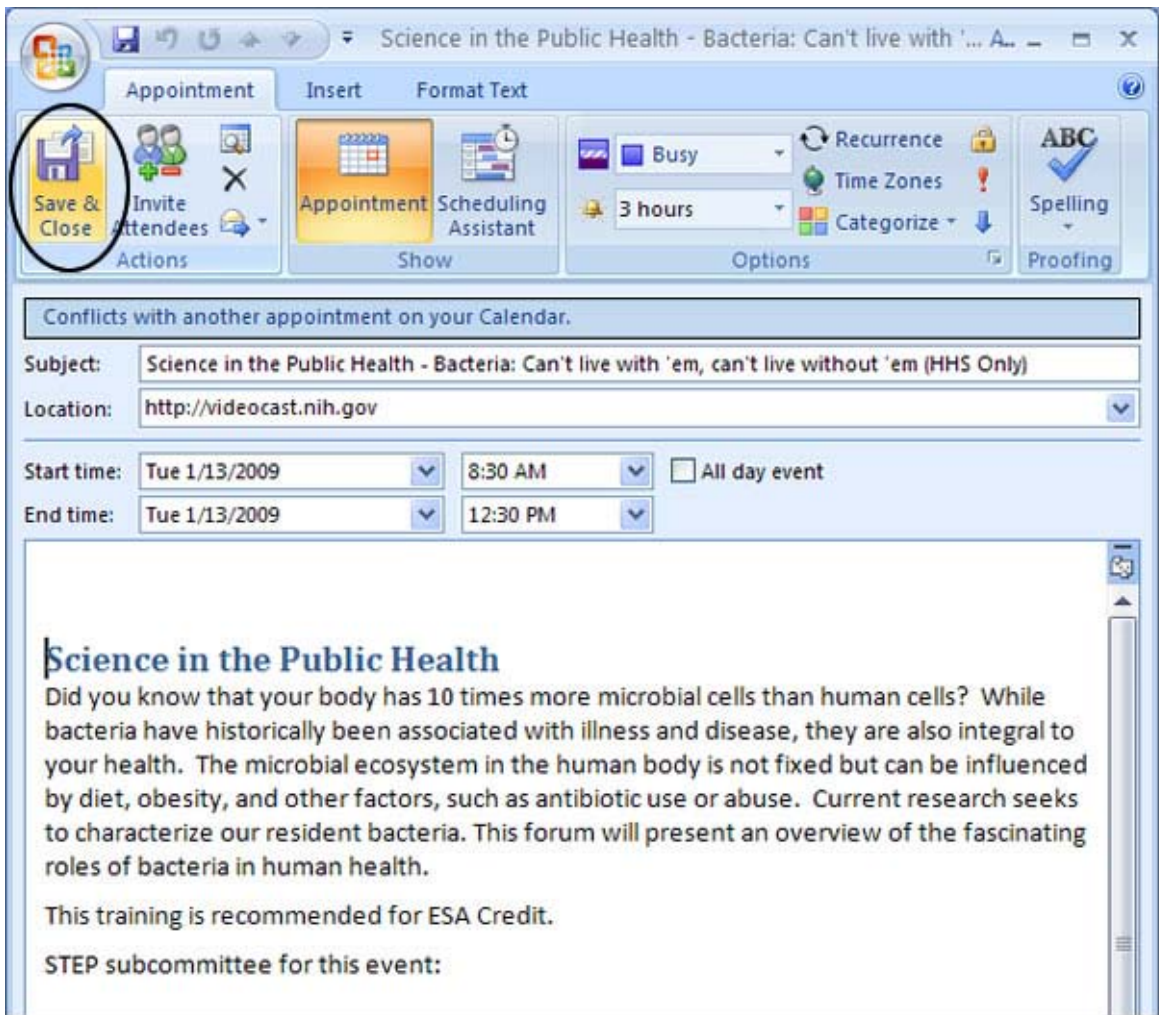
- 1) When you see an Upcoming Event listed on the NIH VideoCast website (<http://videocast.nih.gov>) that is of interest to you, click on its title to view the Summary page where you will find an “iCalendar” button.



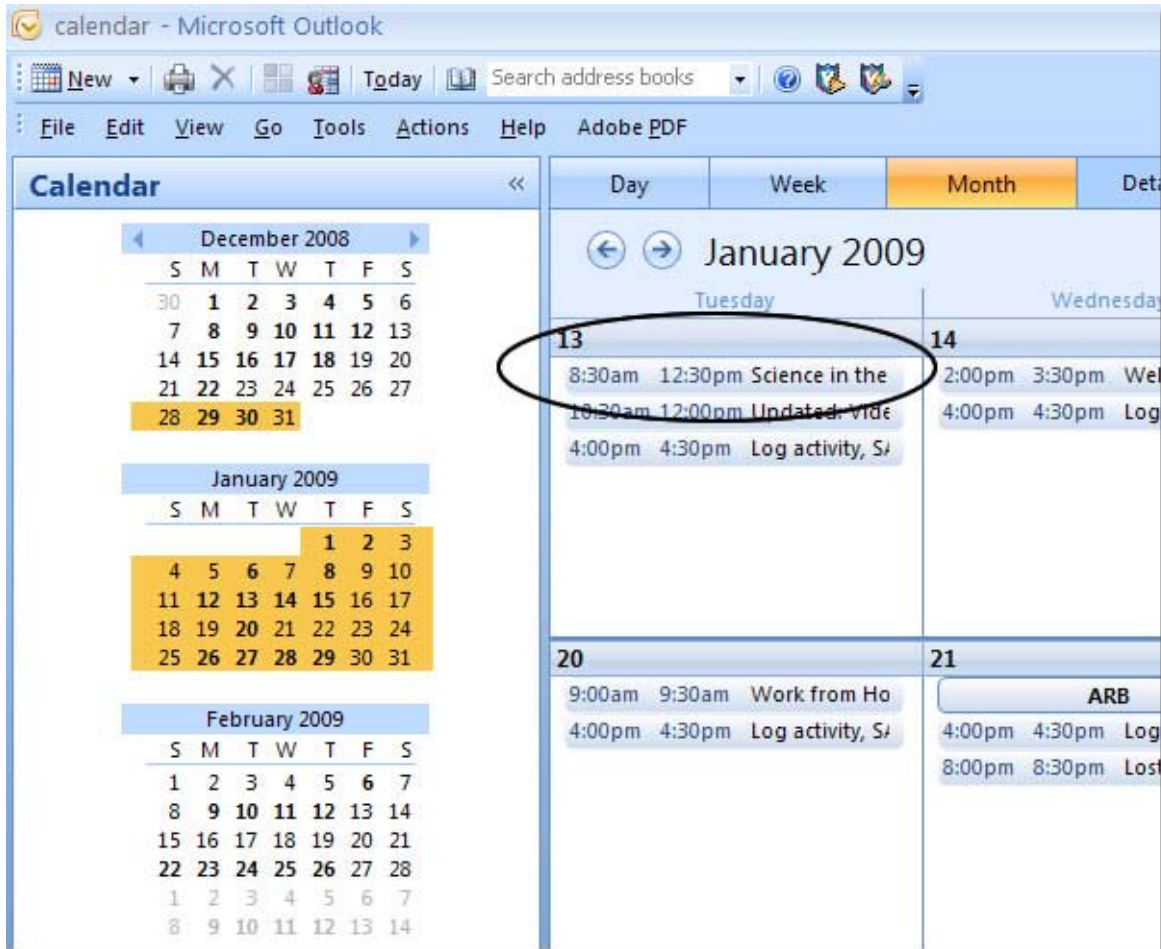
2) Click the iCalendar button, then click "Open" to automatically generate an Outlook Calendar event.



3) Click "Save and Close."



4) And the VideoCast event is automatically added to your Outlook calendar.



TIP: If you still missed that VideoCast, you can search our archives of Past Events to view recorded events at your convenience. Not all VideoCasts are archived, but most will be recorded and made available in the Past Events section within 2 weeks of the live date.

## Questions?

If you have questions about VideoCasts, please contact the NIH Help Desk at <http://ithelpdesk.nih.gov/support> or call 301-496-4357, 301-496-8294 (TTY) or toll free at 866-319-4357 and ask for the NIH VideoCast team.



---

# Going Green by Increasing Printing Efficiency

Going green is good for the environment, good for the NIH Data Center, and good for NIH users. It leads to better health, lower energy costs, and a more sustainable use of our planet.

## NIH greening goals

NIH has established a goal of greening the many activities and business practices occurring daily on its campuses. It is committed to using natural resources responsibly, and that includes the use of energy. The NIH Environmental Management System (<http://www.nems.nih.gov>) sponsors several programs that encourage each organization to use the best available technologies, operations, and practices to increase energy conservation.

The NIH Data Center uses energy primarily in the form of natural gas and electricity to operate, heat, and cool our facilities. As a part of the NIH greening program, the Data Center is working to save power by upgrading to newer Computer Room Air Conditioners and energy-efficient servers, consolidating servers, and using power management tools, as well as high-efficiency power supplies. The Data Center is also exploring avenues for conserving natural resources, including paper, toner, and fuel.

## How you can help

You can be part of creating a greener NIH Data Center by helping us conserve power and resources. Simply using more efficient methods for viewing your Titan output can be one way to do this.

There are several ways to see a Titan output. We will look at three:

- sending it to a continuous form printer
- sending it to a cut-sheet printer
- viewing it online using Interactive Output Facility (IOF)

Of these three, continuous form printers use the most power per page. The cost for maintaining these printers, costs for the special paper and toner that they use, and energy use costs have been steadily increasing and will require a rate increase in the near future.

Therefore, we want to encourage you to consider the less expensive and greener alternatives:

1. Keep output online and view it using IOF; or
2. Send the output to the more energy efficient cut-sheet printers.

By switching to either of these options, you will enable the NIH Data Center to trim the power, paper, toner, and maintenance costs that impact both the environment and the bottom line. Being energy-efficient in handling output will not only decrease power and cooling demands, it will also help us avoid expensive infrastructure upgrades.

---

## Using IOF to save output on disk

To save on your printing costs and reduce the physical space needed for storing and filing printouts, consider using the ISPF Input Output Facility (IOF) to save your job output as disk data sets instead of printing it. You can always print the job from the data set at a later date or edit the data set and then print the output. To avoid printing, you must hold the output using either MSGCLASS=H with SYSOUT=\* or /\* ROUTE OUTPUT HOLD.

Not sure how to save your output? Here's how:

1. On the IOF Job List Menu, select the job whose output you want to save to disk. The example below lists all jobs in the output queue belonging to userid JOEUSER.

```
----- IOF Job List Menu -----( 4 )-----
COMMAND ===>
                                           SCROLL ===> DATA
----- Output Jobs -----
-----JOBNAME--JOBID---ACT-STAT-OWNER----DEST/DEVICE-----RECS-HELD-DAY--TIME
_   1 JOB1      J002621      8 JOEUSER  NIHJES2      1664 037  5:01
_   2 JOB2      J002112      JOEUSER  NIHJES2      5342 037   :13
_   3 JOB3      J001918      JOEUSER  NIHJES2      1787 036  5:01
_   4 JOB4      J001440      JOEUSER  NIHJES2      5342 036  :17
```

2. Let's say that Joe decides that he wants to save JOB2 to a disk data set. To specify JOB2, Joe can enter '2' in the command area or he can tab down and enter an 'S' on the line command next to JOB2 and press "Enter."

```
----- IOF Job List Menu -----( 4 )-----
COMMAND ===> 2
                                           SCROLL ===> DATA
----- Output Jobs -----
-----JOBNAME--JOBID---ACT-STAT-OWNER----DEST/DEVICE-----RECS-HELD-DAY--TIME
_   1 JOB1      J002621      8 JOEUSER  NIHJES2      1664 037  5:01
S   2 JOB2      J002112      JOEUSER  NIHJES2      5342 037   :13
_   3 JOB3      J001918      JOEUSER  NIHJES2      1787 036  5:01
_   4 JOB4      J001440      JOEUSER  NIHJES2      5342 036  :17
```



3. The Job Summary Screen for JOB2 is displayed.

```

----- IOF Job Summary -----
COMMAND ===>                                SCROLL ===> DATA
--JOBNAME--JOBID---STATUS---RAN/RECEIVED-----DAY-----DEST-----
JOB2      J002112  OUTPUT    0:13   2/06/2009 TODAY      NIHJES2
--RC--PGM-----STEP-----PRSTEP---PROC-----COMMENTS-----
0  IEFBR14   STEP1
-----DDNAME---STEP-----PRSTEP---STAT-ACT-C-GRP-D-SIZE-U-DEST-----
_  1  LOG      *                HELD    H   1 H   66 L NIHJES2
_  2  JCL      *                HELD    H   1 H 1226 L NIHJES2
_  3  MESSAGES *                HELD    H   1 H  784 L NIHJES2

```

4. In the command area of the IOF Job Summary screen, Joe enters 'SD' (for Snap Definition) to define the attributes of the data set.

```

----- IOF Job Summary -----
COMMAND ===> SD                                SCROLL ===> DATA
--JOBNAME--JOBID---STATUS---RAN/RECEIVED-----DAY-----DEST-----
JOB2      J002112  OUTPUT    0:13   2/06/2009 TODAY      NIHJES2
--RC--PGM-----STEP-----PRSTEP---PROC-----COMMENTS-----
0  IEFBR14   STEP1
-----DDNAME---STEP-----PRSTEP---STAT-ACT-C-GRP-D-SIZE-U-DEST-----
_  1  LOG      *                HELD    H   1 H   66 L NIHJES2
_  2  JCL      *                HELD    H   1 H 1226 L NIHJES2
_  3  MESSAGES *                HELD    H   1 H  784 L NIHJES2

```

5. The Snap Data Set Options screen appears. The default name of the data set is displayed, using the naming convention *userid.jobname.IOFLIST*; for example, **JOEUSER.JOB1.IOFLIST**. Joe can override the name as well as the UNIT or MGMTCLAS parameters that specify where and how long the data set will be saved.

UNIT and MGMTCLAS are mutually exclusive, so Joe can select only one of these parameters. He can give the parameter he selects one of the following values:

- a. **TEMP** – Output will reside on temporary volume and will be scratched after 7 days online. There are no backups available in this management class.
- b. **NOBACKUP** – Output resides on public volume and is not scratched automatically. No Backups are available.
- c. **DISK2YR** – Output will reside on public volume and will be automatically scratched 2 years after last use. Up to 5 incremental backups are available.
- d. **DISK7YR** – Output will reside on public volume and will be automatically scratched 7 years after last use. Up to 5 incremental backups are available.
- e. **LONGTERM** – Output resides on public volume and is never automatically scratched. Up to 5 incremental backups are available.

In this example, Joe decides to save his output on disk for 7 years using the data set name JOEUSER.COSTREPT.MAY12008:

```

----- IOF Snap Data Set Options -----
COMMAND ===>
                                         Job: JOB1(J002112)

DSNAME          ===> COSTREPT.MAY12008

RECFM           ===> FBA                Record format
LRECL           ===> 133                Logical record length
BLKSIZE         ===>                    Block size
PRIMARY         ===> 30                Primary allocation (in tracks)
SECONDARY       ===> 30                Secondary allocation (in tracks)

UNIT            ===>                    Unit name
VOLUME SER.     ===>                    Volume serial
STORCLAS        ===>                    SMS storage class
MGMTCLAS        ===> DISK7YR            SMS management class
DATACLAS        ===>                    SMS data class

MOD (Y or N)    ===>                    Append data to end of existing data set
PACK(Y or N)    ===>                    Pack data using ISPF method

LABEL           ===>                    Tape data set label type
DATASET NO.     ===>                    Tape data set sequence number

```

- After specifying the output data set parameters, Joe presses ENTER to return to the IOF Job Summary Screen. From this screen Joe can enter SNAP in the command area to write the entire job output to the data set or he can enter 'N' next to each DDNAME he wishes to save.

```

----- IOF Job Summary -----
COMMAND ===>  SNAP                                SCROLL ===> DATA
--JOBNAME--JOBID---STATUS---RAN/RECEIVED-----DAY-----DEST-----
JOB2      J002112  OUTPUT    0:13   2/06/2009 TODAY    NIHJES2
--RC--PGM-----STEP-----PRSTEP---PROC-----COMMENTS-----
  0  IEFBR14   STEP1
-----DDNAME---STEP-----PRSTEP---STAT-ACT-C-GRP-D-SIZE-U-DEST-----
N   1  LOG      *                HELD      H   1  H   66  L  NIHJES2
N   2  JCL      *                HELD      H   1  H 1226 L  NIHJES2
_   3  MESSAGES *                HELD      H   1  H  784 L  NIHJES2

```

- Once Joe has snapped the output, he keeps pressing PF 3 until he returns to the IOF Primary Option Menu. **It is necessary to return to the Primary Option Menu for the data set to be saved.**



---

## Using the cut-sheet printers for output

If you need to print your job, consider routing output to the 2105ES cut-sheet printer instead of the 3900 continuous form printer.

Most job output currently produced on the 3900 continuous form printers are jobs with fewer than 500 pages of output printed in landscape mode. These jobs are ideal candidates for the 2105ES cut-sheet printers.

For jobs that print more than 500 pages of output, we recommend that customers continue to print these jobs on the continuous form printer or save the output to disk.

The following examples show how to route job output to the 2105ES cut-sheet printers. For information about the JCL used, including the FORMS, FORMDEF, PAGEDEF, CHARS, and JESDS statements, see the *Titan Batch Processing* manual at <http://publications.cit.nih.gov/category.asp?category=I007>.

### Printing in landscape mode on both sides of paper with no holes

```
//JOB1      JOB      ( ),MSGCLASS=H
//OUT1 OUTPUT  FORMS=900,FORMDEF=LD00,PAGEDEF=808,DEFAULT=YES,
//          JESDS=ALL,CHARS=GT12
//STEP1 EXEC   PGM=IEBGENER
//SYSUT1 DD   DSN=JOEUSER.SAMPLE.OUTPUT,DISP=SHR
//SYSUT2 DD   SYSOUT=*
//SYSIN DD   DUMMY
//SYSPRINT DD  SYSOUT=*
```

### Printing in landscape mode on both sides of 3-hole punched paper

```
//JOB1      JOB      ( ),MSGCLASS=H
//OUT1 OUTPUT  FORMS=999,FORMDEF=LD99,PAGEDEF=808,DEFAULT=YES,
//          JESDS=ALL,CHARS=GT12,LINECT=59
//STEP1 EXEC   PGM=IEBGENER
//SYSUT1 DD   DSN=JOEUSER.SAMPLE.OUTPUT,DISP=SHR
//SYSUT2 DD   SYSOUT=*
//SYSIN DD   DUMMY
//SYSPRINT DD  SYSOUT=*
```

To avoid printing over the holes on the back of each page, use LINECT=59 to reduce the line count from the default 60 to 59.

---

### Printing in landscape mode on one side of paper with no holes

```
//JOB1      JOB      ( ),MSGCLASS=H
//OUT1 OUTPUT  FORMS=900,FORMDEF=LS00,PAGEDEF=808,DEFAULT=YES,
//          JESDS=ALL,CHARS=GT12
//STEP1 EXEC   PGM=IEBGENER
//SYSUT1 DD    DSN=JOEUSER.SAMPLE.OUTPUT,DISP=SHR
//SYSUT2 DD    SYSOUT=*
//SYSIN DD    DUMMY
//SYSPRINT DD  SYSOUT=*
```

### Printing in landscape mode on one side of 3-hole punched paper

```
//JOB1      JOB      ( ),MSGCLASS=H
//OUT1 OUTPUT  FORMS=999,FORMDEF=LS99,PAGEDEF=808,DEFAULT=YES,
//          JESDS=ALL,CHARS=GT12
//STEP1 EXEC   PGM=IEBGENER
//SYSUT1 DD    DSN=JOEUSER.SAMPLE.OUTPUT,DISP=SHR
//SYSUT2 DD    SYSOUT=*
//SYSIN DD    DUMMY
//SYSPRINT DD  SYSOUT=*
```

## Conclusion

Energy usage is becoming the Data Center's largest single operating cost, at 30 to 40 percent of the operating budget. By reducing Titan's dependence on the 3900 continuous form printer, we hope to work with our customers to become more energy efficient and to increase our IT capacity while maintaining the same energy footprint we have today.

If you need help or have any questions regarding Output Distribution, contact the NIH Help Desk at <http://ithelpdesk.nih.gov/support> or call 301-496-HELP (301-496-4357), 866-319-4357 (toll free), or 301-496-8294 (TTY).



---

# When Was the Last Time You Looked at Titan?

Do you have an image of mainframes – if you think of mainframes at all – as outdated platforms long fallen out of use? Well, it's time to adjust your image and take another look because today's mainframe is not your father's (or mother's) mainframe. In an IT environment of tight budgets, heightened security demands, rising volume, and high energy prices, the new mainframe, aka zServer, is proving its true value. Today's demanding IT requirements make the reliability, cost-effectiveness, security, and scalability of zServer platforms like CIT's Titan an extremely attractive option for data centers and their customers.

So, take another look at the CIT services available on Titan and find out what makes the zServer one of the most durable and adaptive platforms available.

## I. Virtualization

Virtualization has become a really big thing nowadays as data centers struggle with finding space, power, and cooling to deal with more and more servers. Things tend to come full circle – the first and arguably the most robust computer virtualization was accomplished on mainframes many years ago. Through virtualization, customers can save money by reducing hardware purchases.

Titan currently supports two distinct operating systems through z9-Series virtualization technology. The first operating system, z/OS, supports the traditional Titan zServer environment. A second operating system, z/VM, can manage multiple virtual Linux servers on the same z9-Series server. Linux virtual servers can be implemented in a day and can simplify networking requirements and reduce power consumption, resulting in lower costs to our customers.

At a time when customers and data centers are looking for ways to cut energy costs, new applications for the z9-Series server, including using the LAMP application development stack for web and database servers, offer both electrical power savings and the highest level of service. Currently, CIT is implementing a pilot Linux product offering for our customers on the z9-Series enterprise server. This includes a memory upgrade on the z9-Series server to support Linux virtual systems and production implementation of the z/VM operating system and Red Hat Linux.

## II. Security administration

Security remains a major concern of any computing platform, and here zServer platforms have a lot to offer. On the Titan system, **Resource Access Control Facility (RACF)** provides access control and auditing functionality. Not only is RACF known as one of the most mature and scalable security monitors in computing, it also has some interesting features not often found in Windows or Unix environments. It can, for example, set permissions for file patterns – that is, set the permissions even for files that do not yet exist. Those permissions are then used for a file (or other object) created at a later time. In other words, RACF establishes security policies rather than just permission records.

---

RACF has continuously evolved to support such modern security features as digital certificates/public key infrastructure and LDAP interfaces. The underlying z9-Series hardware works closely with RACF. For example, digital certificates are protected within tamper-proof cryptographic processors. Major zServer subsystems, especially DB2 Version 8, use RACF to provide multilevel security (MLS). Some other responsibilities in the RACF area include defining users and passwords, authorizing access privileges to data and other resources to users and groups, performing surveillance reporting, implementing exits, and responding to SAS70 audits.

### III. Automated job scheduling

A real benefit of zServer computing is the ability to run batch jobs — that is to set one program or a series of programs (“jobs”) to run without user intervention, often at idle times when the machine is least busy. This is particularly good for large applications because it allows efficient sharing of resources. A wide assortment of batch applications, including DB2, File Utilities, Batch TSO (REXX) applications, and SAS programs, are executed in the batch environment.

Titan supports several tools to allow automated batch job submission to the z/OS Titan system.

- **CA-7** is a Titan-hosted workload automation product designed to automate critical workloads and IT processes. It supports job submission scheduling, as well as job submission and event triggers on Titan. CA-7 features event-based automation and management, workflow monitoring, enterprise-wide visualization and administration, dynamic critical path management and workload prioritization, embedded workflow, and enterprise reporting.

Customer interface:

- **Web Submit:** This tool allows customers to submit jobs to z/OS Titan for job streams with few dependencies and fairly straightforward scheduling requirements. Just fill out a simple scheduling form at <http://silk.nih.gov/job/scheduler> to specify the JCL and time requirements for the job you wish to submit.
- **Workload Control Center (WCC):** For customers with more sophisticated job scheduling requirements, WCC’s web-based interface provides secure end-to-end visualization, administration, and monitoring of job processing across multiple platforms (such as z/OS and Linux) via a consolidated, web-based console. Using leading-edge portal technology, WCC accesses CA-7 on Titan to schedule job streams on Titan and monitor their execution. CIT’s Automation Support Group will be more than happy to work with you to implement job processing requirements with either Workload Control Center or CA-7.

### IV. Database services

A 2007 Gartner study of information management trends found a marked shift to Linux and z/OS as cost-effective, modern database management system (DBMS) platforms. A closer look at some of the

---

Database Management Systems on Titan, such as DB2 and IMS, should make clear why zServer DBMS are becoming so popular.

## **DB2**

Its reliability and capability in supporting mission-critical business operations make **DB2** for z/OS an excellent tool for large-scale data management. Due to its special environment and the demanding needs of its customers, DB2 for z/OS already offers exclusive features like Multi-Level Security (MLS), extremely large table sizes, and hardware-assisted compression. In addition, DB2 is now starting to acquire Business Intelligence features such as Materialized Query Tables (MQTs) and star schema. While DB2 has always been known for its leading performance and capabilities, **DB2 Release 8** now brings it even closer into alignment with other relational database systems.

DB2 Audio-Image-Video (AIV) extenders are a technology currently being investigated as a CIT service offering. These extenders enable users to take database applications beyond traditional numeric and character data to images, video, audio, and complex documents. Stay tuned for classes and articles on how to take advantage of this technology.

## **IMS**

**Information Management System (IMS)** is a database and transaction management system that was first introduced at NIH to support the old NIH Business System – ADB (which is still in use on Titan until all functions are transferred to the new NIH Business System). Since its introduction, IMS has gone through many changes in adapting to new programming tools and environments.

IMS has two major components, the IMS Database Management System (IMS DB) and the IMS Transaction Management System (IMS TM).

### **IMS DB:**

- Organizes data into a hierarchy – data in each level is dependent on the data in the next higher level;
- Arranges data to ensure integrity and optimize storage and retrieval.

### **IMS TM:**

- Controls I/O (input/output) processing;
- Provides formatting, logging, and recovery of messages;
- Maintains communications security;
- Oversees the scheduling and execution of programs.

IMS has come far from the days when its only programming interface was PL/1. Today, IMS applications and databases can be connected to CICS applications and DB2 databases. Java programs can access IMS databases and services.

---

## V. Web services

Did you know that data stored on Titan can be easily made available on the web? You can get high-performance web access to DB2 and other Titan data using software called **Shadow Web Server**. Shadow Web Server's architecture provides very tight integration into the Titan environment with regard to scalability, security, control, and the protection of zServer assets from risk. Shadow's interface with RACF and use of the "controlled transaction paradigm" allows client-side access to zServer resources while ensuring minimal risk.

Some of the CIT applications that use Shadow Web include Web Sponsor, Chargeback reports, DB2 SQL generator, and RACF administration. Many customers have integrated the Shadow Web services into their applications. Shadow transactions are written in the REXX programming language, a free general-purpose scripting language famous for combining ease of use with power and offering thousands of free add-on tools.

## VI. Printing

Even in a paperless office, sooner or later, everybody has to print something. Titan provides both central and remote print services with a variety of features suited for many different printing needs.

### Central Printing:

- **Infoprint Server:** This Titan feature allows customers to stop tying up their local office printers and print resources (paper and ink) with large print jobs. Instead, simply use Infoprint to send your large print workload from any NIHnet server (i.e. your desktop PC) onto a central Titan print server where we will print, collate, and hole-punch it for you.
- **Print Services Facility (PSF):** PSF is a Titan-based offering for driving Advanced Function Presentation (AFP) on the NIH Data Center's Continuous Form and Cut Sheet Printers. PSF provides a wide variety of data formatting capabilities, including application independent page layouts, formatting of tagged print records, and XML data formatting. Titan PSF helps ensure that your output is printed accurately and completely, without missing or duplicate pages.

### Remote Printing:

- **VPS** is the opposite of Infoprint. It enables you to print any Titan print output generated from such applications as the ADB, DB2, TSO/ISPF, and batch job processing on any printer accessible through NIHnet, such as printers in your local office area.

## VII. Data administration

One of the biggest advantages of our zServer platform is that data on Titan is very secure. Titan automatically backs up all changed data nightly and keeps the last 5 copies of a specific data set version. Titan also performs automatic space-availability and data-availability management between online disk

---

storage, virtual tape storage, and robotic tape devices such as the SL8500. Its operations occur daily at a specified time and allow for keeping a data set for a predetermined period before deleting or moving it. If you accidentally delete your data you can easily recover it yourself with some of the best recovery tools available on the market. In addition to individual data file backups, volume backups are kept both in the Data Center and off-site.

## **VIII. Editors and tools**

The preferred way to perform work on Titan is Time Sharing Option/Interactive System Productivity Facility (TSO/ISPF). This system has an editor and tools to allow users to create, store, print, and delete data sets and to submit and examine jobs.

**TSO/ISPF** is an environment that allows multiple users to access Titan services at the same time. Customers use these facilities to submit jobs and store data sets via the computer's operating system. TSO allows you to execute some programs like SAS and DB2 SQL interactively, while ISPF presents a series of menus from which you may pick the tools you want to use.

### **We are here for you....**

NIH's Titan system draws on decades of innovation and collaboration with our important customers, whose needs in running complex business operations continually shape innovation on our zServer platform. We suggest you look at Titan to reduce cost and complexity and improve security and reliability for your applications.

A Titan z9-Series computer has the equivalent capacity of nearly 150 x86 servers with an 85% smaller footprint and up to 85% lower energy costs. The zServer is a hotbed for business applications – the industry has seen over 600 new applications developed for z9-Series computers in 2007 alone – with momentum in both traditional as well as Linux, Java, and DB2 business intelligence applications.

We are eager to help you and look forward to your feedback and your concerns about Titan – what we do well, what we could do better, and how to better meet your IT requirements. If you have any questions about using this technology, we want to hear from you!

Please visit the NIH Help Desk website (<http://ithelpdesk.nih.gov/support>) or call 301-496-4357, 301-496-8294 (TTY), or toll free at 866-319-4357 to arrange for a Titan consultant to contact you.



---

# Ask the NIH Help Desk

## A solution to over-the-size-limit mailboxes

Is your Outlook mailbox overflowing? Here at the NIH Help Desk, we get quite a few inquiries regarding mailboxes going over the size limits. A very simple solution to this is creating a personal folder and moving items from your mailbox to the personal folder. By storing items in your personal folder, you reduce the size of your mailbox since your personal folder is not part of your mailbox. The instructions below will show you how to create these folders first in Outlook 2003 or, if you have been upgraded to Office 2007, in Outlook 2007.

Users who are migrating from Outlook 2003 to Outlook 2007 shouldn't have any problem with mapping their Personal Folders. Outlook 2007 automatically recognizes the folders so you can access them without any problem.

## How to create a Personal Folder in Outlook 2003

To create a personal folder (also called Archive folder and .PST file) in Outlook 2003, just follow the instructions below.

Step 1 - Click on "File"

Step 2 - Click on "Data File Management"

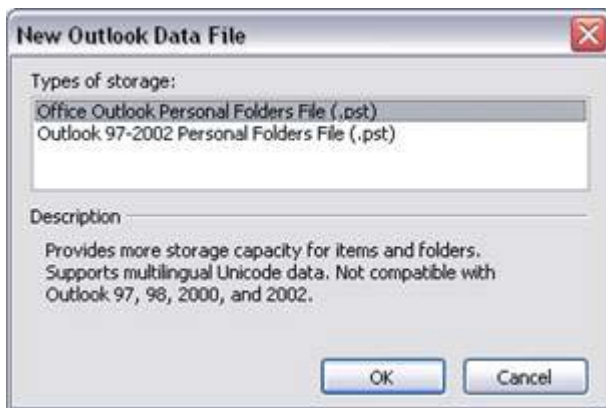
Step 3 - Click on "Add"





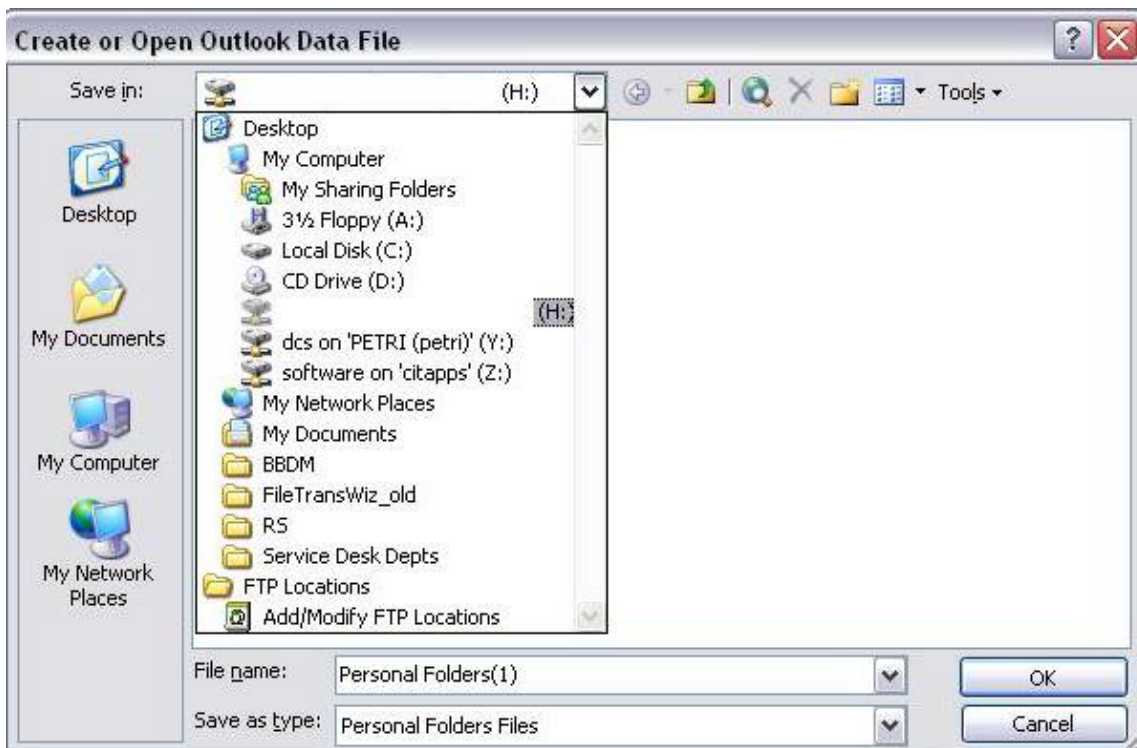
---

This opens a new window showing the “New Outlook Data File” and types of available storage. Select “Office Outlook Personal Folder File (.pst).”



Step 4 – Click “OK”

Step 5 – Select the folder where you would like to save this file (It is always better to save on a Network Drive. A network drive is a folder on a shared drive – not on your own computer – that gets backed up regularly.)



---

Step 6 - Type the desired File Name, then click "OK."

Step 7 - Again type the same name then click "OK."



Step 8 - Close



---

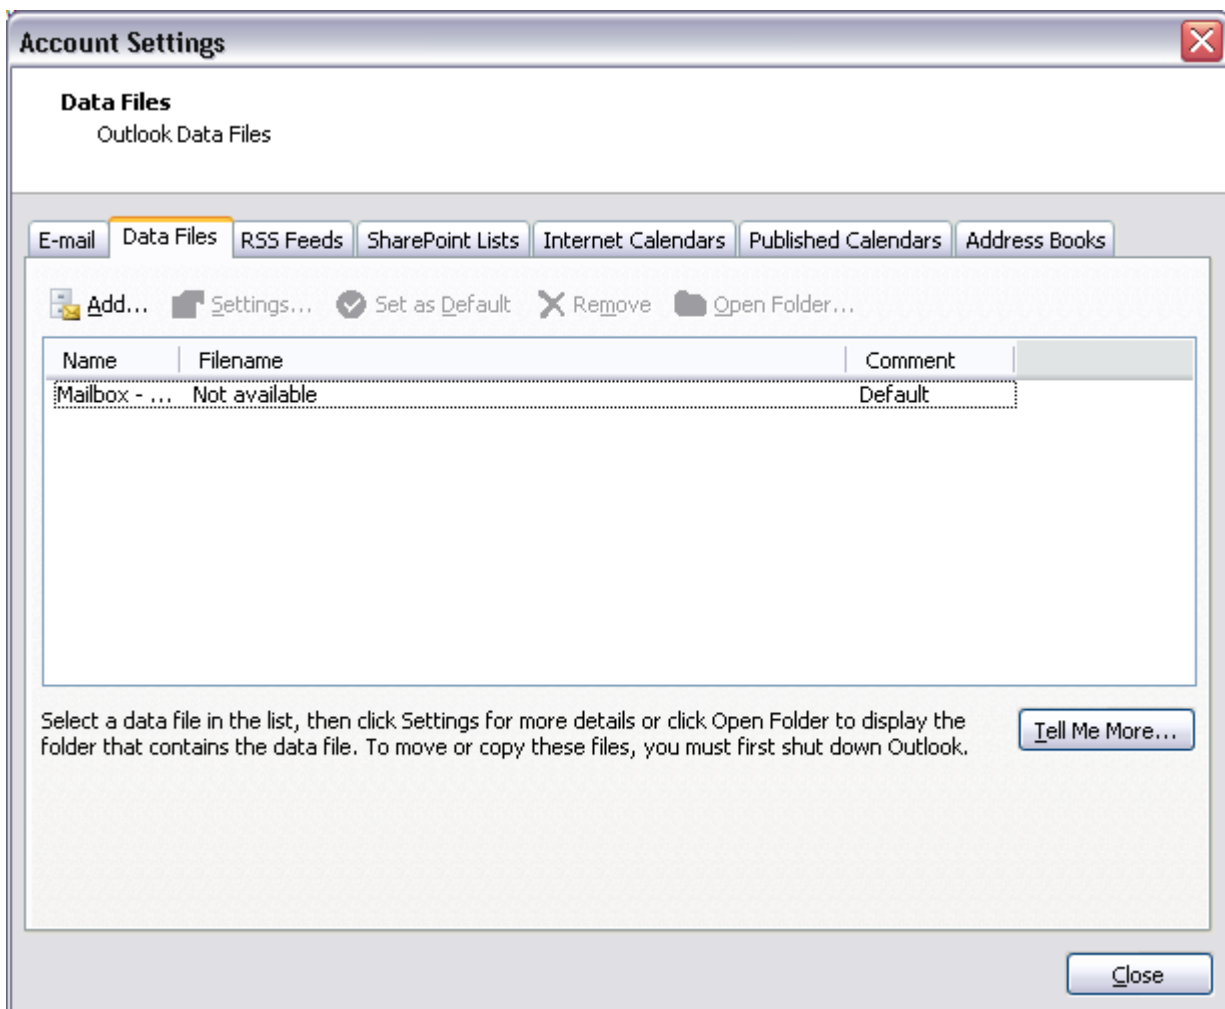
## How to create a Personal Folder in Outlook 2007

To create a personal folder (also called Archive folder and .PST file) in Outlook 2007, just follow the instructions below.

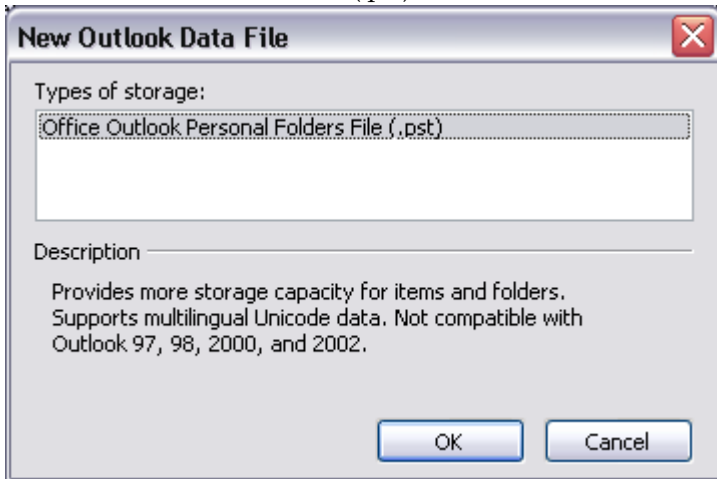
Step 1 - Click on "File"

Step 2 - Click on "Data File Management"

Step 3 - Click on "Add"

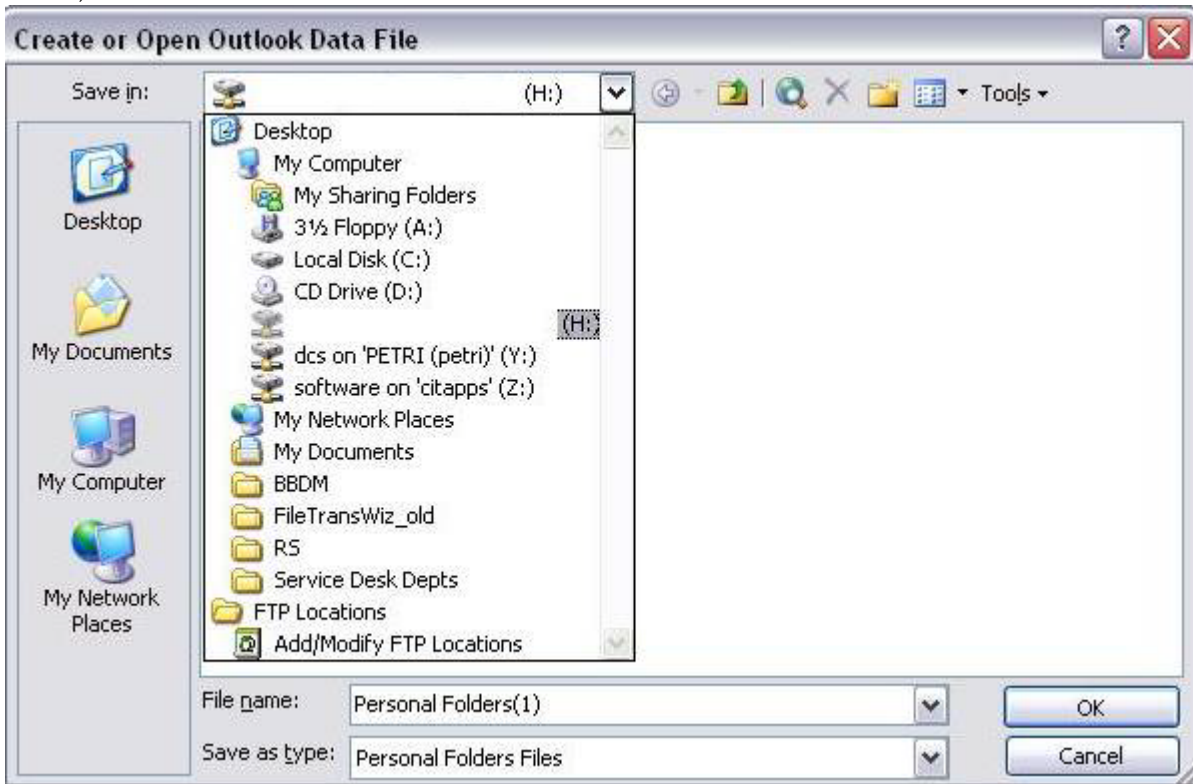


This opens a new window showing the “New Outlook Data File” and types of available storage. “Office Outlook Personal Folder File (.pst)” should be selected.

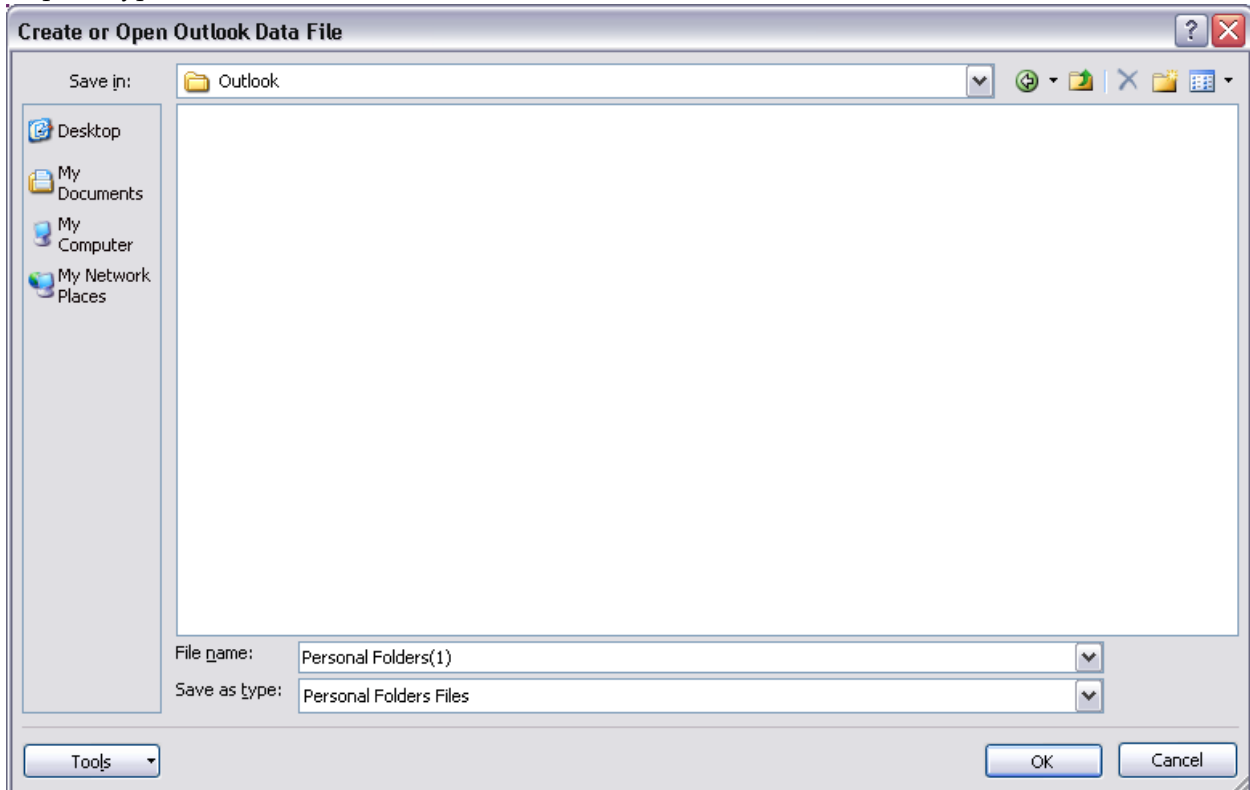


Step 4 - Click “OK”

Step 5 - Select the folder where you would like to save this file (It is always better to save on a Network Drive.)



Step 6 - Type the desired File Name then click "OK."

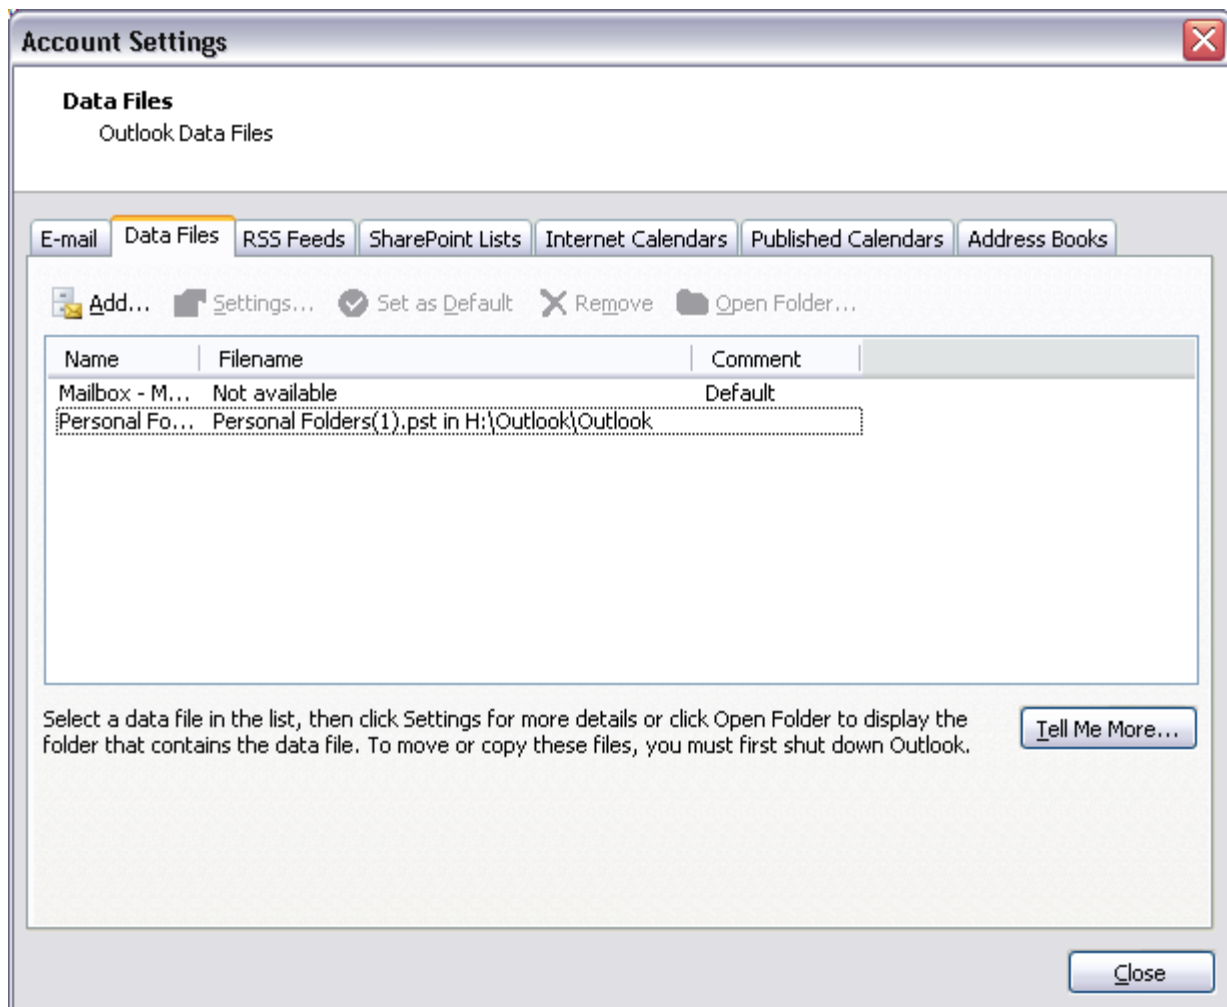


Step 7 - Again type the same name, then click "OK."



---

Step 8 – Close



## Moving items to the personal folder

Once you have created the folder you should see it in your Outlook under "All Mail Folders." To move items to this folder, you can either:

1. Click on the email(s) you want to move and drag the email(s) over to this folder; **or**
2. You can right click on the email(s) and select "Move to Folder." Then you'll see a "Move Items" dialog box from which you select the folder you want to move your emails to, and click "OK."

---

## Questions?

If you have questions or need assistance, please call the NIH Help Desk (<http://ithelpdesk.nih.gov/>) at (301) 496-4357, 301-496-8294 (TTY) or toll free at 866-319-4357.



# Directories and Reference Information

## **NIH Computer Center Hardware and Software**

[<http://cit.nih.gov/ProductsAndServices/ApplicationHosting/RelatedServices/HardwareSoftware.htm>]

## **Computer Services Telephone Directory**

[<http://cit.nih.gov/NR/rdonlyres/CD8200B2-35E6-424C-A1C9-48DA35CE8155/0/TelephoneDirectory.pdf>]

## **Online Services Directory**

[<http://cit.nih.gov/ProductsAndServices/ApplicationHosting/AboutDataCenter/OnlineServices.htm>]

## **Popular Websites for NIH Computer Center Users**

[<http://cit.nih.gov/ProductsAndServices/ApplicationHosting/AboutDataCenter/PopularWebSites.htm>]



## Major Contributors

Darrick Akiyama, DCSS

Laura Bennett, DCSS

Lauren Brush, DCSS

Scott Burge, DCSS

Susan Chacko, DCSS

Phil Day, DCS

John Dussault, DCSS

Judy Fabrikant, DCS

Sarah Fichter, DCSS

Debbie Lang, DCSS

Tim Pickett, DCSS

Norma Stern, DCSS

DCS      Division of Customer Support

DCSS     Division of Computer System Services