

Home

This is the home of the Scientific Computing Services Public (SCSPub) space. This space contains information and guidelines for SLAC users who seek high performance computing and data storage solutions for SLAC research programs and facilities. The Scientific Computing Services (SCS) team provide services and consultation based around centrally-managed, shared resources that can scale beyond individual desktops or workstations. The managed infrastructure is built on a network designed for high-throughput workloads with optimal connectivity to DOE facilities via ESnet. SCS exists to enable and foster all SLAC Science. Our priorities and goals align with the lab Mission, Vision & Values.

Mission

Scientific Computing Services provides storage and computational services that:

- fulfill current requirements and anticipate future needs of its scientific stakeholders;
- are sought after and valued;
- and achieve recognizable efficiencies through shared, common solutions.

Space Index below for search index:

Space Index

Total number of pages: 180

0-9 ... 0	A ... 3	B ... 6	C ... 9	D ... 2	E ... 0
F ... 6	G ... 5	H ... 8	I ... 5	J ... 1	K ... 0
L ... 5	M ... 2	N ... 7	O ... 1	P ... 2	Q ... 0
R ... 6	S ... 26	T ... 3	U ... 12	V ... 3	W ... 0
X ... 0	Y ... 0	Z ... 0	!@#\$... 0		

0-9	A Page: Acquisition Checklist The Acquisition Checklist is intended to help streamline the process of acquiring computing equipment. The first top level bullet lists issues that should be addressed in a technical requirements document. This is a work in progress and feedback on the Page: AFS quota for home directory or group space Most unix accounts still have their home directories in AFS. AFS is also used for some group space. To request an increase in quota or an additional AFS volume, use one of these forms: Self-service for your home directory: https://www.slac.stanford.edu Page: Anonymous FTP at SLAC Anonymous FTP on the central UNIX system at SLAC permits SLAC users and non-SLAC collaborators to exchange files easily. Authorized SLAC UNIX users can store files in FTP space so that collaborators without a SLAC UNIX account may retrieve them. Similarly
------------	--

B Page: Babar LTDA Meeting Notes Page: Babar Meeting Notes Page: Backup and Restore (Unix and AFS) This page has been transferred from the previous website as-is. The information is still relevant. Unix File System Backups at SLAC Unix File Systems For Unix, there are several network file systems managed by the Computing Division, but we will broadly Page: Batch Compute Best Practices and Other Info Specify an output file Use the -o or -oo option to bsub to specify an output file for your batch job. If you do not specify a viable file, the output will be sent via email which when multiplied by the 100s or 1000s can easily overwhelm the mail server. Page: Benchmarking LSI E2660 Storage with RHEL6 Page: Benchmarking MD3460 array with Dynamic Disk Pools The Dell MD3460 controller array can support up to 120 drives with an MD3060e expansion tray attached. Single drive capacities have been steadily increasing, but transfer speeds have remained fairly constant. As a result, RAID6 rebuild times can exceed 24	C Page: CentOS 7 and Chef CentOS 7 is centrally supported at SLAC for the following platforms: VMware virtual machines Bare metal server with devctl for remote console (IPMI / BMC) For desktops or laptops, Ubuntu LTS is the supported choice. Although RHEL 7 is also available if re Page: Chef Configuration Management Table of Contents: Introduction Chef is a configuration management tool (like Puppet, Ansible, SaltStack, CFEngine). It is a tool which manages the configuration of centrally managed Linux servers, compute clusters, and desktops at SLAC. Examples of c Page: Citrix Client for Red Hat As of around May 2018, there were various reports of issues with Citrix at SLAC. The Citrix client wasn't completely broken for all, but it was not ideal for many, and broken completely from some. With the release of a newer firefox, the situation with Page: Colfax/Intel Xeon Phi training slides http://research.colfaxinternational.com/post/2014/10/13/CDT-Slides.aspx http://research.colfaxinternational.com/post/2014/10/13/CDT-Slides.aspx Page: Compute and Clusters Overview The batch system at SLAC uses the IBM Platform Load Sharing Facility, LSF, and is made up of a general farm of batch servers that is open to all SLAC users a rhel6 mpi farm which requires that you request access (send email to unix-admin) and run Page: Compute Cluster Lifecycle <iframe width="1210.3605974395448" height="907.625" seamless frameborder="0" scrolling="no" src="https://docs.google.com/spreadsheets/d/e/2PACX-1vS4JLp4pnJX1PE-QcoH4je0L4PgZlZF31yIUYVPT_r096G-z7GqQJ9KPX7SsxUigy8e8Ob173cYbD2i/pubchart?oid=2043826928&amp;fo Page: Configuring Firefox to enable SPNEGO authentication for Webauth This document describes the process for configuring Firefox to enable SPNEGO authentication. This will allow your browser to use the kerberos tickets you obtained when logging into your linux desktop to access SLAC webauth sites without typing in your pas Page: CPU & Memory limits in LSF with cgrouops Introduction We want to improve the robustness and reliability of the batch system by applying tighter resource controls. The goal is to isolate jobs from each other and prevent them from consuming all the resources on a machine. LSF version 9.1.2 makes u Page: Create/Build a Singularity Container Image The documented method to build a Singularity container image requires using sudo privilege. In this how-to document, we outline how to work with this constraint. Obviously if you have sudo somewhere that you can use to build a Singularity container image,
D Page: Data Transfer with Globus Online Overview Globus Online (GO) is a service that facilitates high throughput data transfer among its endpoints. Many other Universities and Laboratories have endpoints with the GO. SLAC has an public endpoint (slac#osg) that can access most of the NFS, GPF Page: Disk Storage Risk & Lifecycle <iframe width="1206.964864864865" height="720.875" seamless frameborder="0" scrolling="no" src="https://docs.google.com/spreadsheets/d/e/2PACX-1vRQ29MVBv5MPIXgMP6B09zgysxNmeJh7TMNR95OO0K3Jz96A2nRNLdWwRTRzN6OXj877ZX6X9Cty_ES/pubchart?oid=1021695764&amp;for	E

F Page: Fairshare Scheduling Jobs submitted to the general farm of batch systems at SLAC will be scheduled to run according to a cross-queue user-based fairshare priority system. There is more information here www.slac.stanford.edu/comp/unix/package/lsl/currdoc/lsl_admin/index.htm Page: FastX Table of Contents: Quick Start Open your web browser, go to the following URL, authenticate using your SLAC Unix username and password, select either 'Desktop' or 'Terminal'. https://fastx3.slac.stanford.edu:3300 Intr Page: FastX 2 - (deprecated) See the updated instructions for Fastx 3 here: FastX FastX version 2 is no longer supported by the vendor, and will be going away soon. It is left on temporarily to allow everyone to transition to FastX version 3. If you have any problems using FastX ver Page: Fermi meeting notes: 2013-11-21 From: glast-sccs-planning-l@slac.stanford.edu [mailto:glast-sccs-planning-l@slac.stanford.edu] http://mailto:glast-sccs-planning-l@slac.stanford.edu] On Behalf Of Adesanya, Adeyemi Sent: Friday, November 22, 2 Page: Fermi-GLAST Meeting Notes Page: For astore/mstore users	G Page: Get your LSF batch jobs to start faster Specifying a RUNLIMIT You can minimize the time it takes for a general queue job to start running by defining a wall-clock time limit. Instead of explicitly selecting a general queue (short, medium, long, xlong, xxl), just provide the RUNLIMIT argument to Page: Getting Started SLAC User Account A SLAC Unix computer account is required to use our compute and storage services. Staff, users and collaborators affiliated with current SLAC research programs may request a Unix account by contacting their supervisor or research sponsor Page: GPFS What is GPFS IBM General Parallel File System (GPFS) is a high performance parallel filesystem featuring storage virtualization, high availability and is designed to manage large amounts of file data, You can find out more about GPFS in this introductio Page: GPFS storage benchmarks IOZONE tests run locally on single GPFS NSD server fermi-gpfs02 2@md3460 12 NSD 1 MB block system pool 180 disks 128 GB test file (2 x pagepool) /u/sf/jonl/bin/iozone.64bit.linux -i 0 -i 1 -t1 -s 128g -r 1024k Children see throughput for 1 initial wri Page: GPU computing at SLAC: 2013-11-22 2pm From: Deborah Joanne Bard [mailto:djbard@slac.stanford.edu] http://mailto:djbard@slac.stanford.edu] Sent: Friday, November 22, 2013 7:18 PM To: Todd Martinez Cc: Marshall, Stuart L.; Adesanya, Adeyemi; Abel, Tom; Kaehler, Ralf; Brian Moritz; Thomas Peter D
H Page: High Performance Computing at SLAC Overview The batch system at SLAC uses the IBM Platform Load Sharing Facility, LSF, and is made up of a general farm of batch servers that is open to all SLAC users an mpi farm which requires that you request access (send email to unix-admin) and run only Home page: Home This is the home of the Scientific Computing Services Public (SCSPub) space. This space contains information and guidelines for SLAC users who seek high performance computing and data storage solutions for SLAC research programs and facilities. The Scient Page: Home directory in AFS AFS Home Directories: Security Issues SLAC has traditionally encouraged a policy of open exchange of data and programs in its computer systems. However as the Internet has grown and applications have increased in complexity, this policy needs some updatin Page: Hostname or IP address change of CentOS 7 or RHEL 7 host These are the steps for a hostname change or an IP address change for a CentOS 7 or RHEL 7 host IP address change only If the IP address is moving OUT of the current subnet/Vlan, then this needs to be coordinated with Networking. Unless if the host is a V Page: How to blacklist the RHEL 6 Nouveau driver (and install an Nvidia driver) (NOTE: if you follow this procedure, you will introduce a dependency where the kernel and nvidia driver versions must match. This means a centrally-managed system can no longer perform automatic kernel updates without breaking the nvidia driver. Because Page: How to use spack and environment modules to access 3rd party software Describe when someone would need this information. For example "when connecting to wi-fi for the first time". Use the following commands to use a newer version of cmake which is available via spack and modules: \$ bash \$ export MODULEPATH=/afs/slac.stanfor Page: How-to articles Page: How-to change Your Default Unix Shell 1. Login to a RHEL6 SLAC cluster computer (e.g. rhel6-64) Note, even if you are using CentOS7, you should follow these instructions. The update on RHEL6 will propagate to our CentOS7 hosts within a short time. ssh <username>@rhel6-64.slac.s	I Page: Index Page: Installing Citrix Receiver on CentOS 7 Download Citrix Receiver for Linux https://www.citrix.com/downloads/citrix-receiver/linux/receiver-for-linux-latest.html https://www.citrix.com/downloads/citrix-receiver/linux/receiver-for-linux-latest.html As of 2019-Dec-13, this was the tar ball I downl Page: Installing YFS on Ubuntu Desktop NOTE: SLAC IT does NOT support AFS on desktops. AFS is being retired. To access AFS on Ubuntu desktops (to migrate data, for example), use the "File" application, "Other Locations", "sftp://centos7.slac.stanford.edu". screenshot1.pngscreenshot2.pngscre Page: Intel Developer Tools and Libraries The Intel Parallel Studio XE Composer Edition (C/C++/Fortran) is available to all SLAC users. Our license restricts the number of concurrent builds. There are no license restrictions on the runtime libraries. https://software.intel.com/en-us/intel-paralle Page: Interactive Login Pools Monthly Reboots The following interactive login pools are rebooted on the first Sunday of each month, staggered between 4:00 AM - 4:30 AM Pacific Time. fastx3 https://confluence.slac.stanford.edu/display/SCSPub/FastX NoMachine https://confluence.slac.stanford.edu/display/SCSPub/FastX

J Page: Jupyter at SLAC Jupyter is a web based analysis and coding environment. It supports multiple different programming languages, but is mostly centered around python development. The main advantage over standard IDEs is that it provide immediate code execution and inline gr	K
L Page: LCLS Meeting Notes Page: LCLS Unix account password process LCLS URAWI admins need the following information from the RES database: - is the account DISABLED (or ENABLED) - is the password EXPIRED PROPOSED SOLUTION: 1) Create 2 new fields in the RES (Resource Enumeration System) database to track two different att Page: Linux Docker containers with LSF Docker containers may provide a 'lightweight' solution for running multiple linux environments on a single host. Science collaborations could create Docker 'images' that encapsulate their libraries and executables. These images could be portable across mu Page: Linux Server Monthly Reboots The morning of the first Wednesday of the month is designated as a maintenance reboot window for some Linux servers. Linux servers can be automatically rebooted by taylor or chef, or they can be manually rebooted by a unix-admin team member. This monthly Page: LSF and RHEL 7 We have a a couple of RHEL 7 hosts available to LSF for testing. To submit a job, bsub -q rhel7 <Job> Please note that RHEL 7 is still being developed and is not yet a production environment at SLAC, so you will find that there may be some missing pieces	M Page: Managing and Deploying Applications on OpenStack Thursday March 19, Noon - 1 PM SLAC Building 52, Room 103 (Mad River Conference Room) openstack.jpg Come listen to Vish Ishaya, OpenStack veteran and expert, talk about private cloud computing! Topics discussed will include strategies for deploying applic Page: Meeting Notes Our meeting notes with other groups.
N Page: Nagios at SLAC Summary https://nagios.slac.stanford.edu/ https://nagios.slac.stanford.edu/ What is Nagios? Nagios http://www.nagios.org/documentation is an open-source monitoring tool. It is used at SLAC to automatically watch key hosts and services, and to contact appr Page: Nagios Central Service Level Objectives - 2011-09-08 From Shirley: Notes from the Nagios meeting on 2011/09/08 Attending: Tony Johnson, Charlotte Hee, Tom Glanzman, Richard Dubois, Shirley Gruber, Yemi Adesanya, John Bartelt The objective for centralized Nagios support is that there will be less work for al Page: Nebula Nebula private cloud computing, OpenStack, Chris Kemp July 2014. Powerpoint slides: Chris C. Kemp - SLAC.pptx Page: New Hardware Page: News and Announcements 11-April-2014 - noric and yakut aliases being removed on 5-May-2014 On 5-May-2014, the noric and yakut aliases will be removed. Please use the following names to access the compute interactive login machines: rhel5-32 rhel5-64 rhel6-32 rhel6-64 It was pre Page: NoMachine S3DF NoMachine information: If you are using the SLAC Shared Scientific Data Facility (S3DF), please see this page for information about the S3DF NoMachine service: https://s3df.slac.stanford.edu/public/doc/#reference?id=nomachine https://s3df.slac.stanford.edu/public/doc/#reference?id=nomachine Page: nvidia-automatic-builds-via-dkms Automatic build and install of the Nvidia kernel module/driver using Dynamic Kernel Module Support (DKMS). Red Hat based systems ship with an nvidia-compatible graphics kernel module, and user space X11 driver) called nouveau https://nouveau.freedesktop.org	O Page: OpenNebula talk 10-July-2014 Bringing Private Cloud Computing to HPC and Science - SLAC - July 2014 .pdf
P Page: Parallel Computing Overview All SLAC users can run parallel jobs on the "bullet" shared cluster. It has 5024 cores. Each cluster node is configured as follows: RHEL6 64bit OS x86 nodes 2.2GHz Sandy Bridge CPUs 16 cores per node 64GB RAM per node QDR (40Gb) Infiniband for Page: PPA Lustre filesystem 2014 upgrade Here are some benchmarks for PPA Lustre filesystem that was upgraded in March 2014. System specs: Lustre server version 2.5.1 on RHEL6.5 1 MDS and 8 OSS servers using Dell R610 systems Four LSI Engenio (Dell MD3260) arrays with dual-redundant controllers	Q
R	S

Page: [Red Hat backporting FAQ](#)

Explanation: <https://access.redhat.com/security/updates/backporting> <https://access.redhat.com/security/updates/backporting> More info below (from <https://access.redhat.com/solutions/57665>) What is backporting and how does it affect Red Hat Enterprise L

Page: [Red Hat Software Collections](#)

Red Hat Software Collections <https://www.softwarecollections.org/en/> "Software Collections give you the power to build, install, and use multiple versions of software on the same system, without affecting system-wide installed packages" For instance, RHEL

Page: [Remote Access](#)

Page: [Restoring files using TSM](#)

[This page has been transferred from the previous website mostly as-is. The information is still relevant.] To restore your own backup files, you generally need to be on the same machine used to originally backup the files. The TSM (IBM Tivoli Storage Ma

Page: [Restoring Your Mail Spool File](#)

This page only applies to individuals who use the Unix Mail Spool system for mail delivery. It does not apply to anyone who uses Exchange/O365. Your Unix mail spool file from /var/spool/mail may be restored from ITSM backups. However, before you attempt

Page: [RHEL 7](#)

11-April-2014 Red Hat's RHEL7 High Touch Beta Program ended last month. The Release Candidate for RHEL 7 is expected very soon. After that, GA (official release) is expected (no official dates, but we guess before the end of May 2014). SCS will offer an

Page: [Samba Unix Storage Access](#)

Samba (SMB/CIFS protocol) allows you mount remote SLAC Unix storage on your local desktop or laptop if you are on the SLAC network. Authentication is done using your SLAC Windows Active Directory username and password. CentOS 7 installation: `sudo yum`

Page: [Scientific Computing Services - Mission Statement](#)

Scientific Computing Services provides storage and computational services that: fulfill current requirements and anticipate future needs of its scientific stakeholders; are sought after and valued; and achieve recognizable efficiencies through shared, com

Page: [SCS Town Hall for Unix Community April 23, 2015](#)

SCSTownHall.2015.04.23.pptx

Page: [SCS Town Hall for Unix Community August 6, 2015](#)

SCSTownHall.2015.08.06.pptx

Page: [SCS Town Hall for Unix Community January 14, 2016](#)

SCSTownHall.2016.01.14.pptx SCSTownHall.2016.01.14.pdf

Page: [SCS Town Hall for Unix Community January 22, 2015](#)

SCSTownHall.2015.01.22.pptx

Page: [SCS Town Hall for Unix Community July 22, 2014](#)

SCSTownHall 2014_07_22.pptx

Page: [SCS Town Hall for Unix Community March 2, 2017](#)

SCSTownHall.2017.3.2.pdf NERSC_Gerber.2017.3.2.pdf NERSC_Fagnan.2017.3.2.pdf

Page: [SCS Town Hall for Unix Community May 12, 2016](#)

SCSTownHall.2016.05.12.pdf

Page: [SCS Town Hall for Unix Community October 14, 2014](#)

SCSTownHall.2014.10.14.pptx

Page: [SCS Town Hall for Unix Community September 22, 2016](#)

Scientific Computing Services: SCSTownHall.2016.09.22.pdf PIV-I: UnixTH_20190922_PIV-I.pptx Firewall Tightening Update: 2016-09-22 UNIX town hall - cybersec update.pdf

Page: [SCS Town Hall for Unix Community September 28th, 2017](#)

SCSTownHall.2017.9.28.pdf

Page: [SCS Town Hall for Unix Users 2013-12-12](#)

SCSTownHall 2013 12.pptx

Page: [SCSTown Hall for Unix Community 10-Apr-2014](#)

SCSTownHall.2014.04.10.pptx SCSTownHall.2014.04.10.pdf

Page: [SLAC Compute and Storage Resources](#)

30,000 cores, ~300TFlops/s 150 GPU, ~2PFlops/s 35PB disk 60PB on tape 100Gbps internal network 2x100Gbps external network connectivity to ESnet 10Gbps backup network

Page: [Slurm Batch](#)

Slurm is a batch scheduler that enables users (you!) to submit long (or even short) compute 'jobs' to our compute clusters. It will queue up jobs such that the (limited) resources compute resources available are fairly shared and distributed for all users

Page: [Software](#)

Page: [SSH](#)

Table of Contents: Projects: SSH Inbound Connections Reduction https://slacprod.servicenowservices.com/kb_view.do?sysparm_article=KB0012232 * SLAC IT Cyber Security Owns this project, for more information please see the link. (SLAC Active Directory Login

Page: [SSH and Shared Service Accounts](#)

(Copied from an old web page. Needs clean up.) <http://www.slac.stanford.edu/icon/blank.gif> SSH and Shared Accounts Previously SLAC used a locally customized version of SSH that supported forwarding AFS tokens during login. Unfortunately, the latest versio

Page: [SSHFS Unix Storage Access](#)

SSHFS allows you mount remote SLAC Unix storage onto your local desktop or laptop. You can use SSHFS from anywhere (eg, home or remote network). Authentication is done using your SLAC Unix username and password. SSHFS uses the SFTP protocol and SSH auth

Page: [Stakeholder priority on the Shared Farm](#)

10-26-2021: As we migrating to SDF, decommissioning old hardware and RHEL6, we will no longer actively update the fair shares in this page - some of the major stakeholders no longer use LSF batch system in large scale. The Shared (General) Farm consists o

Page: [Status and Announcements](#)

Live Status More detailed metrics and monitoring can be found at: Grafana <https://grafana.slac.stanford.edu/?orgId=1> Nagios <https://nagios.slac.stanford.edu> Ganglia <http://ganglia.slac.stanford.edu:8080/> PlatformRTM <https://farmrtmweb.slac.sta>

Page: [Storage](#)

	This page is a work in progress. Space for your scientific data or for output from your research or analysis is available in your Unix home directory (a relatively limited amount) or in shared storage space. Your experiment might have its own shared stor Page: Storage as a Service (StaaS) [2022-07-11: StaaS will be superseded by S3DF beginning in FY23. Please refer to https://sdf.slac.stanford.edu/public/doc https://sdf.slac.stanford.edu/public/doc/#/ for more information.] Description: Storage as a Service (StaaS) is a SLAC shared file s Page: Storage benchmarks Unix Storage Project link. Page: System Overview Computing Interactive Computing Batch Computing GPU Computing Storage Local Scratch AFS GPFS Software Module LSF
T Page: Technical Overview of new SDF Storage, March 23rd 2020 PowerPoint Slides Video recording of zoom presentation https://stanford.zoom.us/rec/play/tJlklUGsrD43HNTBuASDAf95W461Kq2shikd-KdZxU7gB3dXZgClbuAWN-WgdHhGJdo9V0SW38VgcX4s?continueMode=true Page: Thunderbird and Owl Email Configuration Introduction IMAP email access at SLAC is limited to the internal SLAC network. If you are offsite, you need to connect to VPN to use the IMAP protocol. Alternatively, you can use and configure the Owl plugin for Thunderbird. Owl allows Thunderbird to use Page: Transferring Data	U Page: Ubuntu Desktop How-To SUIT and SCS Instructions for installing a Chef centrally managed Ubuntu linux desktop. Ubuntu installation instructions: Install Ubuntu 16.04 or 18.04 Chef installation instructions: Run the following command to bootstrap Chef central configuration manag Page: Ubuntu Desktop testing progress There are two Ubuntu 16.04 Desktop test boxes being used by SUIT: Jacob Demo PC86881 IP: 134.79.68.12 HWADDR: 78:2b:cb:b3:c2:a2 Franklin Pham PC89158 IPADDR: 134.79.68.110 HWADDR d4:be:d9:2f: 0e:2a Both the above desktops have been installed with Ubuntu 1 Page: Ubuntu System Administration System Administration tips for Ubuntu Ubuntu Security Information Tracker CVE Database: http://people.canonical.com/~ubuntu-security/cve/ CVE Tracker: https://launchpad.net/ubuntu-cve-tracker https://launc Page: Ubuntu/CentOS 7 Desktop Scope of Support Plans for modern linux desktop support. Desktop Linux Distributions supported at SLAC Recommended Linux Distribution: Ubuntu Long Term Support 18.04 or 20.04 Long term support (LTS) releases are for 5 years. 18.04 = YY.MM of release date (released April Page: UbuntuDesktop Introduction Ubuntu LTS is the recommended Desktop platform at SLAC. It is centrally managed by SLAC IT Help Desk and Office of the CIO (OCIO) Unix Platform computing. Chef is used for configuration management and compliance. To request a centrally mana Page: UNIX Disk Space Costs It is reasonable to ask for up to a total of 10-20 GB with little or no justification. Beyond that we generally would like to see some sort of explanation of what you need the space for. When there are larger quantities of data involved, it can become imp Page: Unix Town Hall, August 27th, 2020 UnixTownHall.2020.08.27.pdf CyberUnixProject-UnixTownHall08272020.pdf NERSC_Update_August_2020.pdf Zoom videoconference recording: https://stanford.zoom.us/rec/share/65BYd6HM2VpJWpHG6kzwBpEaRlv0eaa81SdNqPsPnx02ALAYxEFWGL6Gn5IQOccb?startTime=159854671000 Page: Unix Town Hall, February 7th 2019 UnixTownHall.2019.02.07.pdf Page: Unix Town Hall, June 28th 2018 UnixTownHall.2018.06.28.pdf Page: Unix Town Hall, November 14th, 2019 UnixTownHall.2019.11.14.pdf NERSC-9.pdf NERSC-IAM.pdf Zoom videoconference recording: https://stanford.zoom.us/recording/share/89JZGe6BxaBQIn_G0C07tMkZ-r-ZEB-UBTI5UBKqFuwlumekTziMw https://stanford.zoom.us/recording/share/89JZGe6BxaBQIn_G0C07tMkZ-r-Z Page: User Documentation Page: Using Environment Modules This page describes high level usage of SLAC's implementation of environment modules. environment modules are a way of dynamically loading arbitrary programs into your unix shell environment. examples include the ability to switch between different versi

V Page: Vagrant and VirtualBox You can use Vagrant and VirtualBox to quickly bring up Virtual Machines on your Mac, Windows, or Linux desktop. Download Vagrant https://www.vagrantup.com/ https://www.vagrantup.com/ Download VirtualBox https://www.virtualbox.org/ https://www.virtualbox.org/ Page: Viewing HTML email attachments with alpine for alpine to view html attachments: 1) edit ~/.mailcap, remove any html lines, add this line: text/html; elinks -dump %s; nametemplate=%s.html; copiousoutput 2) when viewing the email in alpine: select ">" for "ViewAttch" press return to view html select Page: VNC on Unix The use of VNC for remote connections is not recommended by the Cyber Security team and VNC is not a centrally-supported service. The supported solution for remote graphical X11 connections for Unix is FastX. https://confluence.slac.stanford.edu/display/S	W
X	Y
Z	!@#\$