

PBS/Cylc Integration Project @ BoM

David Block • Senior Director • Altair Engineering



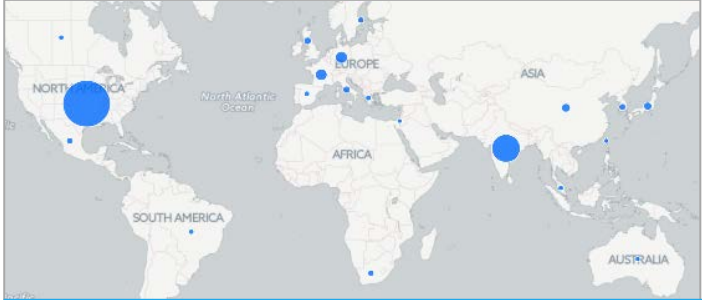
Altair

ALTAIR ENGINEERING

A bit about Altair Engineering



ALTAIR AT A GLANCE



Founded 1985
Headquartered in Troy, MI US

69 offices
in 24 countries

\$333M
2017 Revenue

50+
ISV partners under our unique,
patented licensing model



2000+
Engineers, scientists and creative thinkers

5000+
Customer installations globally

60,000+
Users



PBS PROFESSIONAL

PBS PROFESSIONAL® IS ALTAIR'S INDUSTRY-LEADING
WORKLOAD MANAGER AND JOB SCHEDULER FOR
HIGH-PERFORMANCE COMPUTING (HPC) ENVIRONMENTS.



MARKET-LEADING SECURITY

EAL3+ Security & Cross-Domain Support (aka MLS)



PBS Professional is the **first and only workload manager** to achieve Common Criteria EAL3+ Certification

PBS Professional is the **first and only workload manager** to support multi-level security via Red Hat SELinux support



RHEL MLS configurations are Intelligence Community Directive 503 certified and used in operations

Commercial | Hardened | Certified Secure



Industry-leading workload manager and job scheduler for high-performance computing

[Downloads](#)

[Support Options](#)

[Communities and Commercial](#)

[Source Code](#)

[Fork Me on GitHub](#)

What is PBS Professional[®]?

PBS Professional software optimizes job scheduling and workload management in high-performance computing (HPC) environments – clusters, clouds, and supercomputers – improving system efficiency and people’s productivity. Built by HPC people for HPC people, PBS Professional is fast, scalable, secure, and resilient, and supports all modern infrastructure, middleware, and applications.





Search or jump to...

Pull requests Issues Marketplace Explore



PBSPro / pbspro

Unwatch 66 Star 195 Fork 151

Code Pull requests 15 Projects 0 Insights

An HPC workload manager and job scheduler for desktops, clusters, and clouds. <http://www.pbspro.org>

hpc pbs pbs-professional job-scheduler cloud cluster green provisioning hooks plugins scheduling

663 commits 3 branches 6 releases 56 contributors

Branch: master New pull request Create new file Upload files Find file Clone or download

visheshh	analyze_server_log does not return any values if only start time is s...	Latest commit 5b3c404 6 days ago
.appveyor	Add simple test in AppVeyor	14 days ago
.github	Added a section for log output in the template	5 months ago
.travis	Add sanitize build support in Travis	7 months ago
buildutils	PP-1174: Fix copyright headers in several files	9 months ago
doc	Man page updates	3 months ago
docker	Build PBS Pro on OpenSUSE Tumbleweed	12 days ago
m4	LIBICAL_API2 macro is defined when libical version 1 is passed for build	5 days ago
src	LIBICAL_API2 macro is defined when libical version 1 is passed for build	5 days ago
test	analyze_server_log does not return any values if only start time is s...	4 days ago
win_configure	Add simple test in AppVeyor	14 days ago



PLAGIARISM

The following slides have largely been taken from a presentation prepared by:

Jacinta Richardson, Software Engineer, Model Build Team, BoM
jacinta.richardson@bom.gov.au

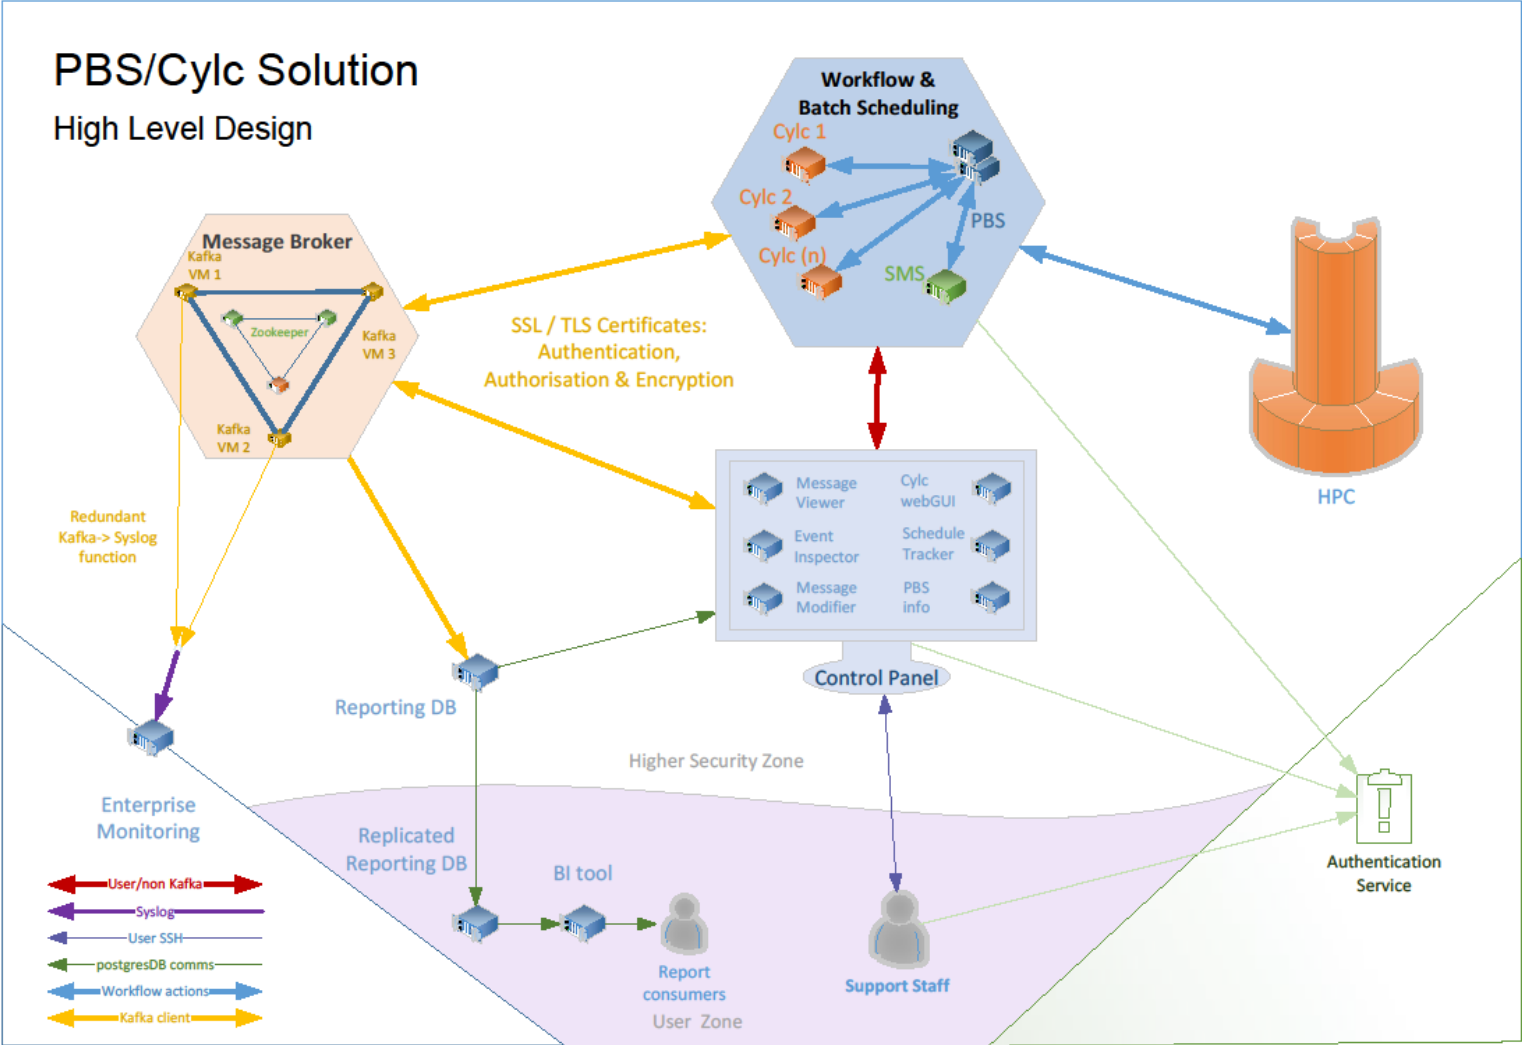


THE BOM TRANSITION

- Still using SMS for the majority of our suites
- Chose to move to cylc+rose
- Cylc+rose delivered to production December 2017
- First suite running in production using cylc January 2018
- Project still continues... this talk briefly details why



PLANNED ARCHITECTURE



ON-GOING GOALS

- Integration with SMS until models are fully ported
- Reduce inter-suite coupling:
 - Current SMS suites are tightly coupled
 - New cylc suites are polling for specific data
 - Want event-based triggering
- Future proofing (rich data schema → Business Intelligence)
- On a suite error, send (meaningful) message to enterprise monitoring system
- Reporting and analysis of work-flow and errors
- Easy to develop with, debug with, and support
- Security and assurance

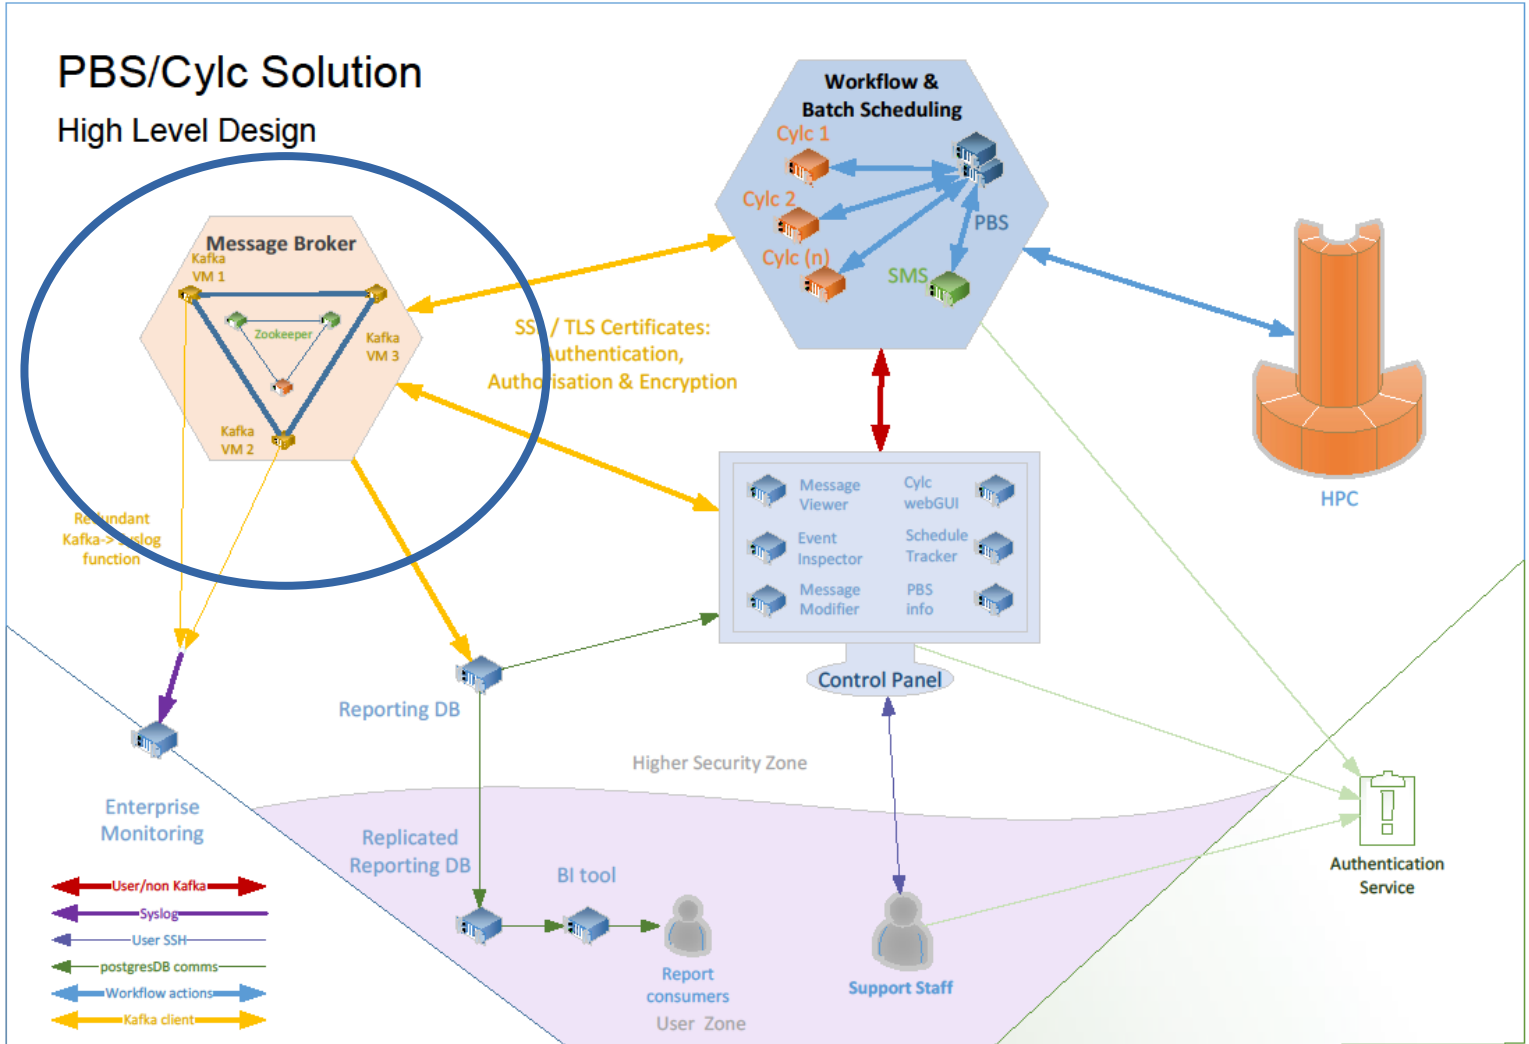


BENEFITS TO PARTNERS

- US Navy is also a big SMS user, interested in cylc
 - Tim Whitcomb came to our cylc workshop
 - Security is a big focus of theirs too
- Hilary is our key consultant
 - All improvements are in turn available to all cylc users
- Altair providing enterprise support for cylc, PBS and the wider tool-chain (that we're working together on creating)



KAFKA CLUSTER



BENEFITS OF KAFKA

- Kafka is a message broker
- Kafka is an Apache open source project used by the New York Times, Pinterest, airbnb, Cloudflare, coursera, cisco, Etsy, Uber, LinkedIn, Loggly, Netflix, Oracle, PayPal, Spotify...
- Reduces inter-suite coupling
 - Producers don't need to know about their consumers
 - Consumers don't need to know about their producers' internals
- Each suite can have its own topic to publish to, but read from many



FUTURE PROOFING

- Kafka gives us independence from:
 - Suites changing their internal structure
 - Architectural changes (e.g. changing enterprise monitoring tool)
 - System architectural changes (e.g. reprovisioning more cylc VMs and moving suites to load-balance them)
- Kafka allows us to change (or add) future schedulers
- Our protocol buffer schema gives us independence from Kafka

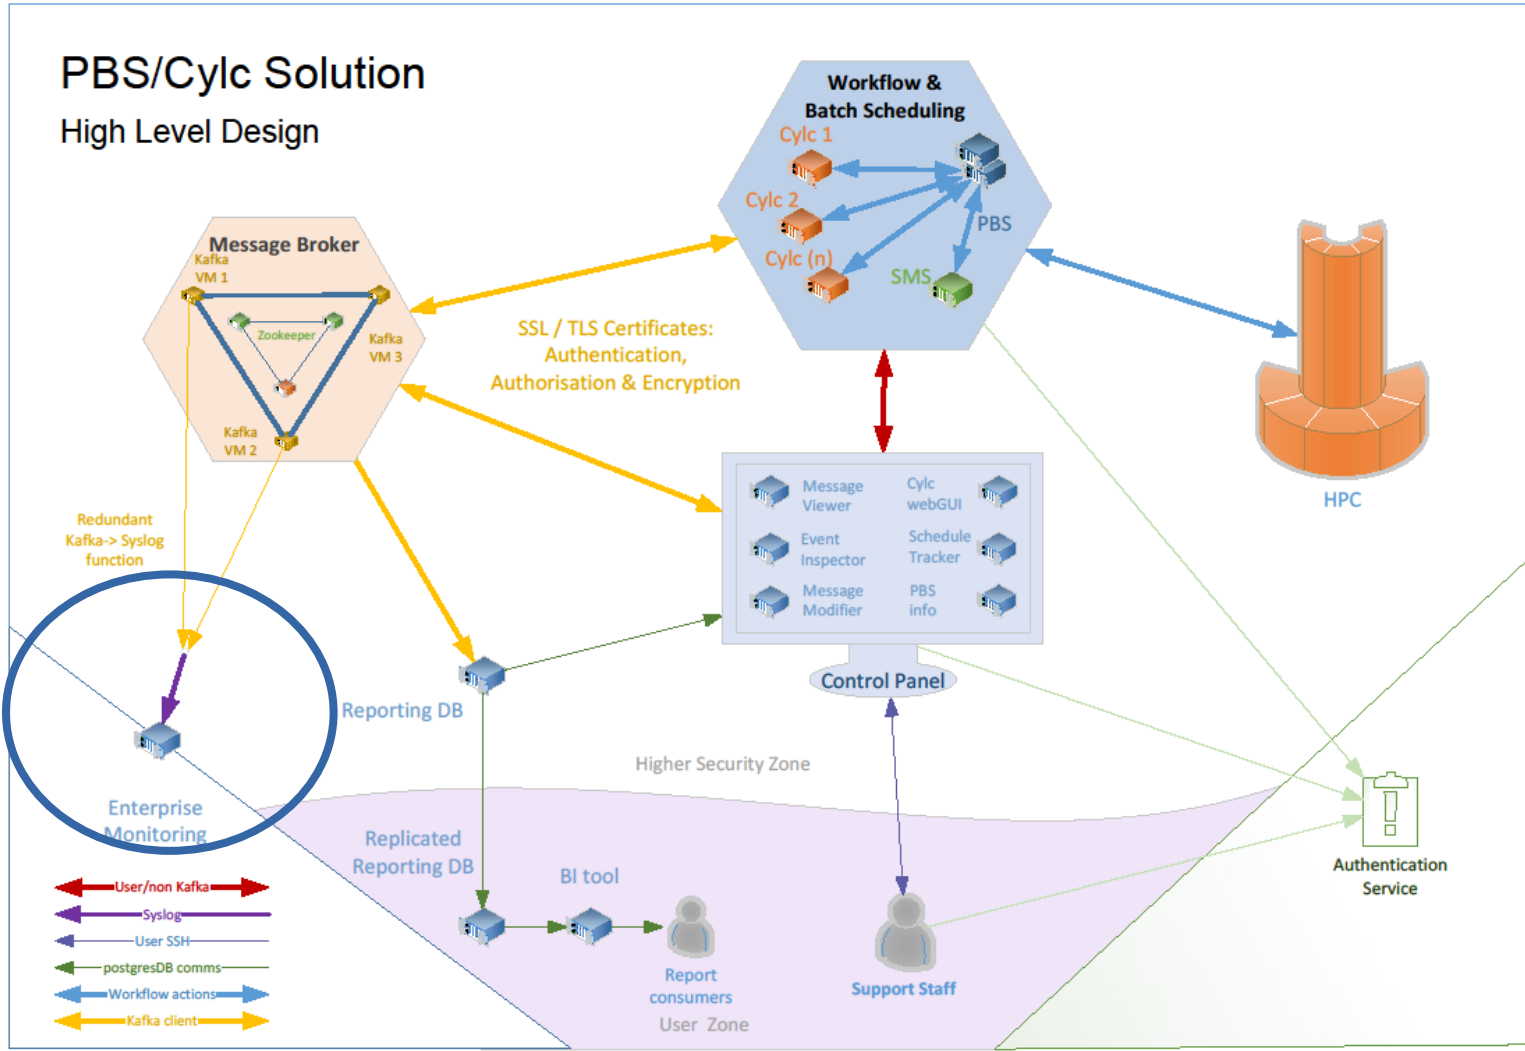


ON-GOING GOALS

- Integration with SMS until models are fully ported ✓
- Reduce inter-suite coupling:
 - Current SMS suites are tightly coupled ✓
 - New cylc suites are polling for specific data ✓
 - Want event-based triggering ✓
- Future proofing (rich data schema → Business Intelligence) ✓
- On a suite error, send (meaningful) message to enterprise monitoring system
- Reporting and analysis of work-flow and errors
- Easy to develop with, debug with, and support
- Security and assurance



ERRORS GO TO ENTERPRISE MONITORING

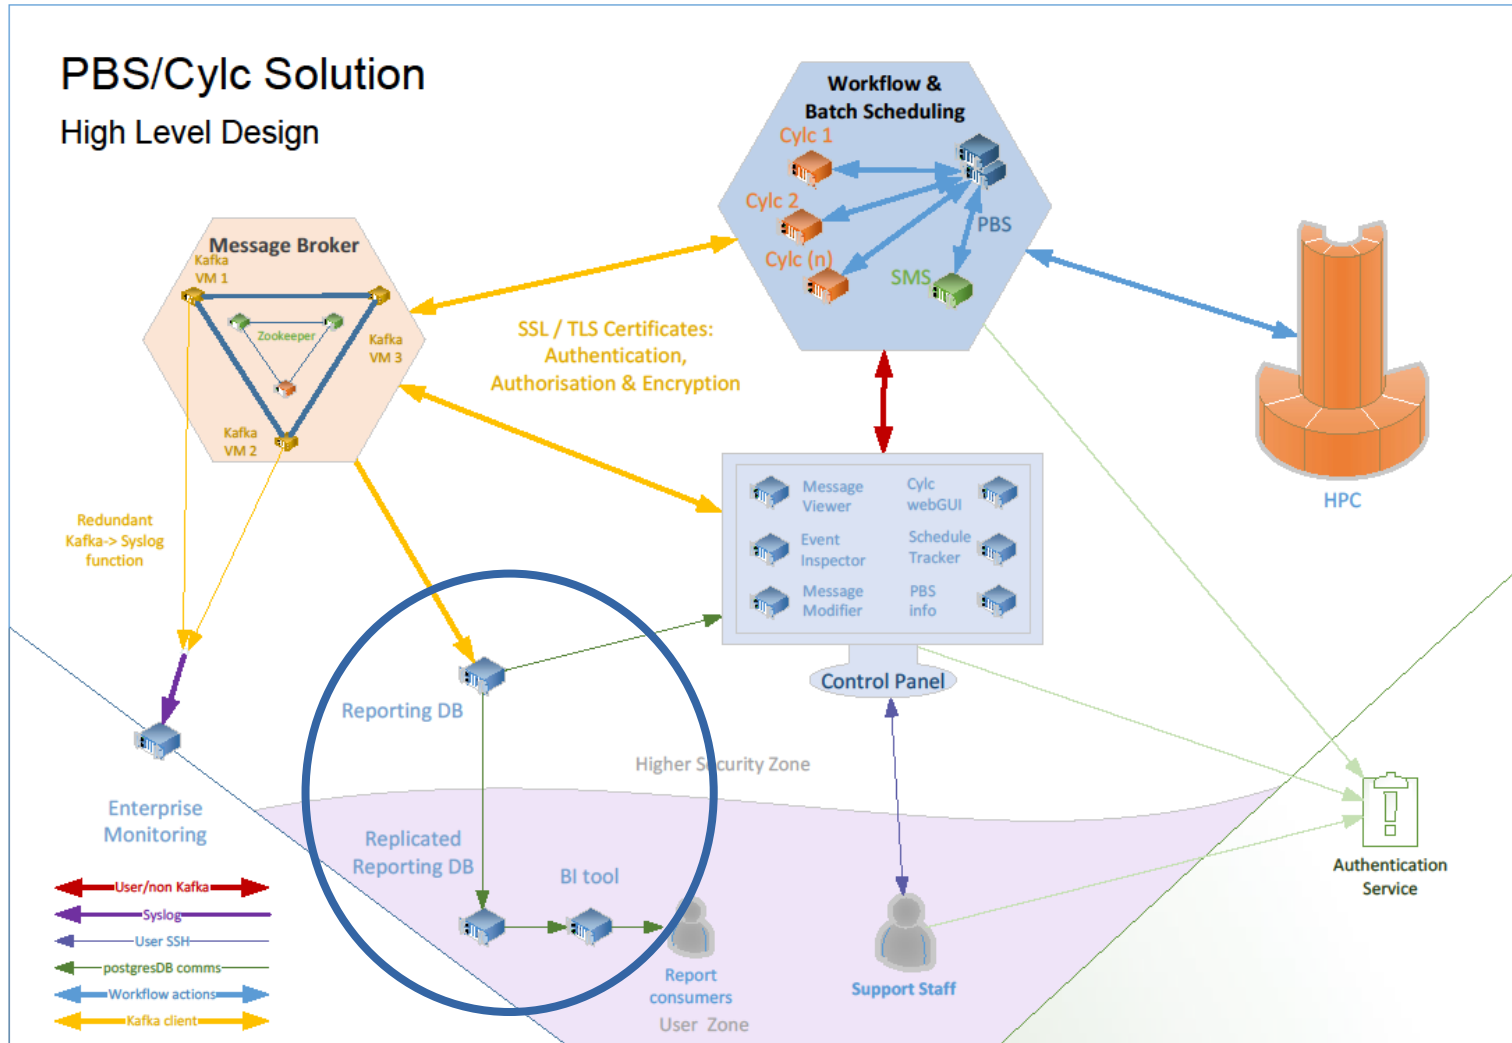


ON-GOING GOALS

- Integration with SMS until models are fully ported ✓
- Reduce inter-suite coupling:
 - Current SMS suites are tightly coupled ✓
 - New cylc suites are polling for specific data ✓
 - Want event-based triggering ✓
- Future proofing (rich data schema → Business Intelligence) ✓
- On a suite error, send (meaningful) message to enterprise monitoring system ✓
- Reporting and analysis of work-flow and errors
- Easy to develop with, debug with, and support
- Security and assurance



REPORTING DATABASE



REPORTING DATABASE

- Collects logs from cylc daemon and state database
- Collects logs from PBS
- Collects events from cylc suites
- Replicated to user zone for wider use
- Explorable, but we'll also have pre-generated reports

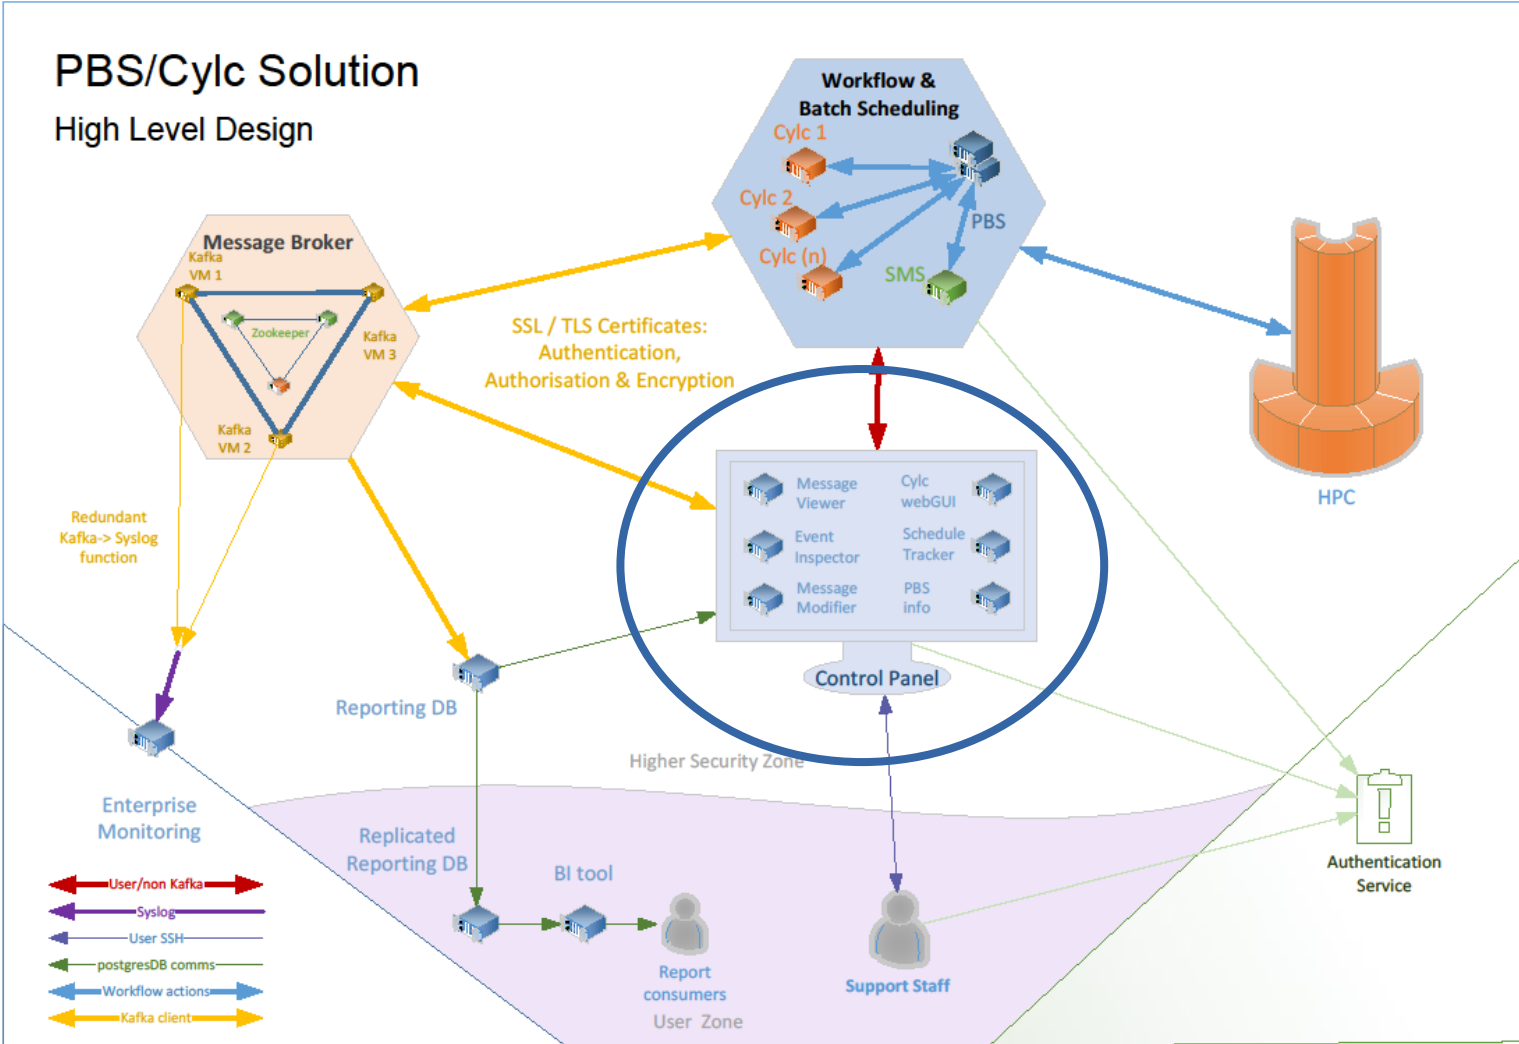


ON-GOING GOALS

- Integration with SMS until models are fully ported ✓
- Reduce inter-suite coupling:
 - Current SMS suites are tightly coupled ✓
 - New cylc suites are polling for specific data ✓
 - Want event-based triggering ✓
- Future proofing (rich data schema → Business Intelligence) ✓
- On a suite error, send (meaningful) message to enterprise monitoring system ✓
- Reporting and analysis of work-flow and errors ✓
- Easy to develop with, debug with, and support
- Security and assurance



EXTRA TOOLING



EXTRA TOOLING

- Cylc GUIs to run in browser
- Kafka message fixer
- Event database viewer
- Kafka topic streaming viewer
- Scheduler prediction tool



ON-GOING GOALS

- Integration with SMS until models are fully ported ✓
- Reduce inter-suite coupling:
 - Current SMS suites are tightly coupled ✓
 - New cylc suites are polling for specific data ✓
 - Want event-based triggering ✓
- Future proofing (rich data schema → Business Intelligence) ✓
- On a suite error, send (meaningful) message to enterprise monitoring system ✓
- Reporting and analysis of work-flow and errors ✓
- Easy to develop with, debug with, and support ✓
- Security and assurance



SECURITY

- Cylc security assessment and report
- Penetration testing
- Analysis of cylc's ssh use, exploration of ssh-free cylc
- PAM authentication plugin for cylc to remove on-disk passphrase and certs
- Kafka's support for topic authentication and authorisation
- Integration of web tools with standard PAM



ON-GOING GOALS

- Integration with SMS until models are fully ported ✓
- Reduce inter-suite coupling:
 - Current SMS suites are tightly coupled ✓
 - New cylc suites are polling for specific data ✓
 - Want event-based triggering ✓
- Future proofing (rich data schema → Business Intelligence) ✓
- On a suite error, send (meaningful) message to enterprise monitoring system ✓
- Reporting and analysis of work-flow and errors ✓
- Easy to develop with, debug with, and support ✓
- Security and assurance ✓



Thank you...

Please ask Hilary your questions. :)

