

Investigating Potential Reliability and Robustness of Standards-Based Grid Computing Systems

Chris Dabrowski
National Institute of Standards and Technology

Presented at GGF15 in Boston, USA
October 6, 2005

Motivation

Vision: Future global information infrastructure will rely on emerging standards for Web Services and Open Grid Services Architecture


Question #1: – Will future distributed systems designed in conformance to Web Services and Grid standards achieve levels of robustness, scalability, and performance required for critical enterprise applications?

Question #2: – As industrial grid systems grow in size, can unplanned interactions among distributed components lead to emergence of undesirable patterns of incoherent and chaotic behavior?

Question #3: – Can we identify areas in GGF specifications that might lead to in implementation of operational Grids that are unreliable or that experience unexpected failures?

Possibility of emergent behaviors

- **Possible Concerns:** System designs may lead to interactions under failure conditions that result in emergent behaviors and unexpected performance degradations. The scaling of grid systems may, in and of itself, result in emergent behavior that adversely affects system behavior.



Definitions of Terms

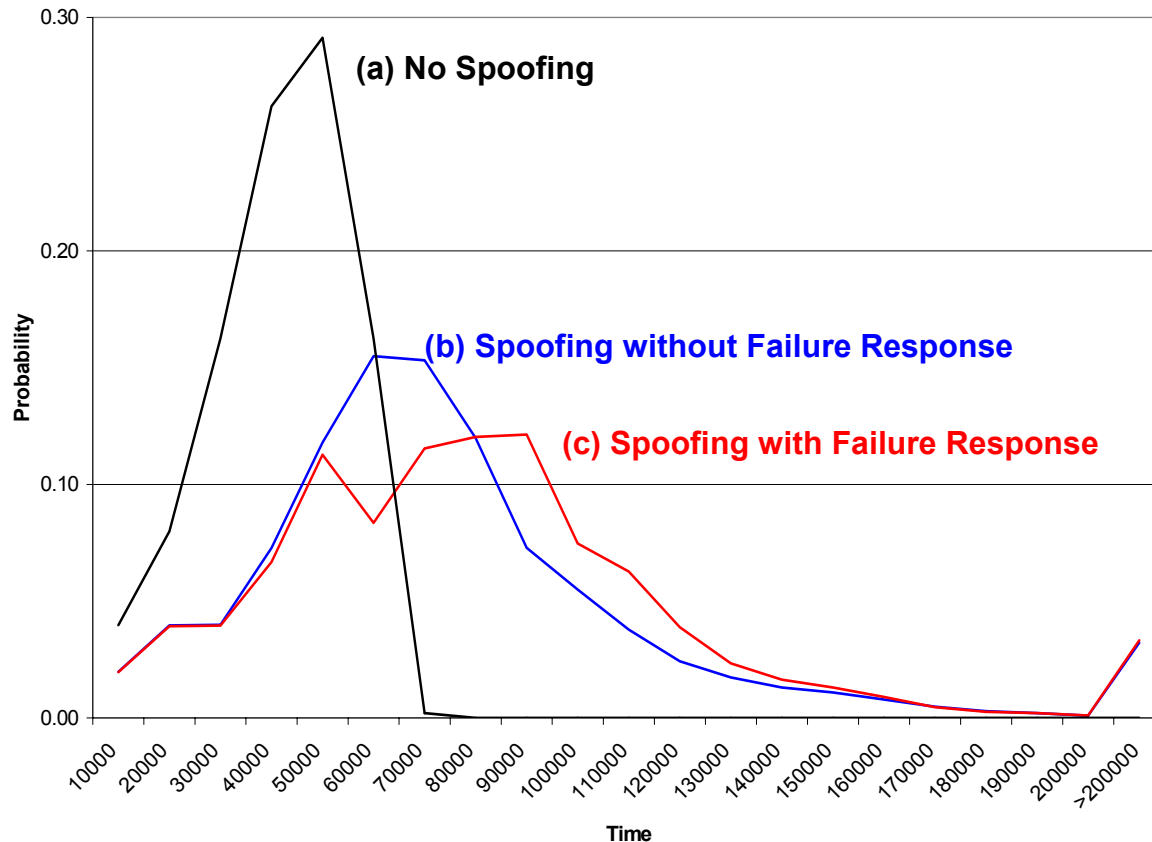
Reliability (system)

- “An attribute of any system that consistently produces the same results, preferably meeting or exceeding its specifications. *Free On-line Dictionary of Computing*”, <http://foldoc.doc.ic.ac.uk/foldoc>
- “In engineering in general, reliability is the capacity of a component or a system of such components to perform as designed”, *Wikipedia*, <http://en.wikipedia.org/wiki>

Robustness

- “... an ability to recover gracefully from the whole range of exceptional inputs and situations in a given environment”, *Free On-line Dictionary of Computing*
- “In terms of computer software, refers to the resilience of the system when under stress - for example, when running a large number of processes, or when starved of memory or storage space, or when confronted with an application that has bugs or is behaving in an illegal fashion....”], *Wikipedia*

Example: Performance Degradation Caused by Spoofing Activity in Simulated Grid



Comparative Probability Density Functions (PDFs) for Application Completion Times given:
(a) No Spoofing, (b) Spoofing without Failure Response, and (c) Spoofing with Failure Response.

Next Steps?

What actions should GGF take to assess potential of robustness and reliability of grid systems developed with GGF specifications?

- Form RG to serve as focus of research that addresses questions of grid reliability and robustness? **[This group forms recommendations]**
- Initial meeting of this BOF (GGF12 Brussels) resulted in draft RG charter. Can we restart this effort?

Research Group for Reliability and Robustness?

- **Motivation:** *(from previous slides)* As grid systems are increasingly commercialized and grow in size, they are likely to be subjected to volatile and uncertain conditions that endanger or severely degrade their effectiveness in everyday use.
- **Question to be addressed:** How can we determine that the web-service and grid standards currently being developed will enable large-scale grids to detect and overcome failures to provide a level of reliability and robustness needed for industrial and scientific purposes?
- **RG Focus/Purpose:**
 - Identify issues related to reliability and robustness in grid computing systems designed in conformance to Web Services and Grid standards
 - Make recommendations, and explore methods, for improving reliability and robustness of standards-based grid systems developed for critical enterprise applications and production systems.

Research Group Scope

- Identify reliability concerns in industrial and scientific grid systems developed on the basis of Web Services and Grid Standard Specifications.
 - especially, those that can be directly (or indirectly) related to behaviors required or implied by Web Services and Grid Standard specifications
- Develop recommendations/guidelines for
 - Specifications that enhance reliability of deployed systems
 - Standards-based grid systems that exhibit high reliability
- Identify (or foster research that identifies) test methods and metrics for evaluating grid systems reliability, including the ability to detect, and respond to various kinds of failures, such as failures of individual components, links, as well as entire subnetworks. .
- Explore (or foster research that explores) algorithms or protocols for improving grid reliability and robustness.
- Examine how mechanism for ensuring reliability (reliable grid ftp, monitoring service, replication service, checkpointing, etc.) interact with grid specifications/standards. Exploring relationship between mechanisms and standards
- Define minimum performance levels/thresholds for grid system reliability

Research Group Goals

- To organize forums that allow researchers, application developers, and others to present results of their work on grid reliability and related work and to exchange information. The RG may provide information, through web-pages, mailing lists, white papers, best practices documents, and other publications, on critical reliability issues as they relate to:
 - GGF standard specifications and those of related bodies
 - Defined or deployed grid services and software
 - Real-world Grid system applications
- To promote and facilitate:
 - Collaborations between researchers in grid systems reliability
 - Access to testbeds and simulation models that investigate reliability issues
- Represent the general experiences and findings on grid systems reliability of application developers using Grid technologies

Possible Products/Deliverables

- Documents: requirements, best practices, and case studies, reports identifying specifications with reliability concerns
 - best practices document: year 1 study on impacts of reliability issues on GGF work (reliability requirements for standards, identification of impacts within GGF standards (which standards are impacted, what changes might be made), empirical studies on reliability in grid systems (including statistical studies), & case studies
- Workshops that collect requirements and case studies; serve as focus to bring community together.

Web Services

- [1] A. Szalay, et al., “Web Services for the Virtual Observatory”, MSR-TR-2002-85, Microsoft Research, August 2002
- [2] L. Melloul, T. Pering, and A. Fox, “WeSCoS: A First Step Toward Web Programming,” an unpublished draft, Stanford University
- [3] Y. Daoi, et al., “Managing Web server performance with AutoTune agents”, IBM SYSTEMS JOURNAL, VOL 42, NO 1, 2003, pp. 136-149
- [4] A. Karp, “E-speak E-xplained”, HP Labs Technical Report, HPL-2000-101 August 7, 2000
- [5] Charu Aggarwal, Joel L. Wolf, Philip S. Yu, “Caching on the World Wide Web”, Knowledge and Data Engineering, 1999
- [6] B. Urganonkar, et al. “Maintaining Mutual Consistency for Cached Web Objects”, Proceedings of the 21st International Conference on Distributed Computing Systems (ICDCS-21), Phoenix, Arizona, April 2001
- [7] P. Chandra, et al. “Darwin: Resource management for value-added customizable network service”, *Proceedings of the 6th IEEE International Conference on Network Protocols (ICNP '98)*, 1998
- [8] A-C Huang and P. Steenkiste, “Hierarchically-Synthesized Network Services”, unpublished draft, Department of Computer Science, Carnegie-Mellon, 2003
- [9] F. Curbera, I. Silva-Lepe and S. Weerawarana, “On the Integration of Heterogeneous Web Service Partners”, Proceedings of the Workshop on Object-Oriented Web Services (OOPSLA '2001)
- [10] E. Kiciman, L. Melloul, and A. Fox, “Position Summary: Toward Zero-Code Service Composition”, a position paper from some Stanford researchers
- [11] V. Machiraju, J. Rolia, A. van Moorsel, “Quality of Business Driven Service Composition and Utility Computing”, HP Labs Technical Report 66, 2002

Web Services (cont.)

- [12] J. Li, M. Yarvis, and P. Reiher, "Securing Distributed Adaptation", *Computer Networks*, Vol. 38, No. 3, 2002, pp. 347-371
- [13] V. Tosic, D. Mennie, and B. Pagurek, "Dynamic Service Composition and Its Applicability to E-Business Software – The ICARIS Experience", *Proceedings of the WOOBS (Workshop on Object-Oriented Business Solutions) workshop* (at ECOOP 2001), Budapest, Hungary, June 18, 2001, pp. 95-108
- [14] S. Chandrasekaran and S. Madden and M. Ionescu, "Ninja Paths: An Architecture for Composing Services over Wide Area Networks", CS262 class project writeup, UC Berkeley (2000)
- [15] P. Felber, et al. Failure Detectors as First Class Objects, *Proceedings of the International Symposium on Distributed Objects and Applications (DOA'99)*, IEEE Computer Society Press, September 5-7, 1999, p. 132
- [16] C. Marchetti, A. Virgillito, and R. Baldoni, "Design of an Interoperable FT-CORBA Compliant Infrastructure," *Proceedings of the European Research Seminar on Advances in Distributed Systems (ERSADS)*, 2001
- [17] S. Frolund et al. "Building Dependable Internet Services with E-speak", *Proceedings of the Workshop on Dependability of IP Applications, Platforms, and Networks*, June 26, 2000, held in conjunction with the 2000 International Conference on Dependable Systems and Networks, IEEE Computer Society, and also available as Hewlett-Packard Labs Technical Report 2000-78
- [18] D. Liang et al., "A Fault-Tolerant Object Service on CORBA," *The Journal of Systems and Software*, vol. 48, 1996
- [19] M. Chen, E. Kiciman, E. Fratkin, A. Fox, E. Brewer, "Pinpoint: Problem Determination in Large, Dynamic Internet Services", *Proceedings of 2002 International Conference on Dependable Systems and Networks (DSN)*, IPDS track, Washington, DC, June 23-26, 2002

Grid Computing

- [20] S. Tuecke, et al., Open Grid Services Infrastructure (OGSI), Version 1.0, Global Grid Forum, June 23, 2003
- [21] M. Baker, R. Buyya, and D. Laforenza, “Grids and Grid technologies for wide-area distributed computing”, *Software Practice and Experience*, 2002
- [22] The Grid Report, The Commercial Implications of the Convergence of Grid Computing, Web Services, and Self-managing Systems, Bloor Research – North America, August 2002
- [23] Z. Juhasz, A. Andics, and S. Pota, “Towards A Robust And Fault-Tolerant Discovery Architecture For Global Computing Grids”, in Distributed and Parallel Systems – Cluster and Grid Computing, which contains the Proceedings of the 4th Austrian-Hungarian Workshop on Distributed and Parallel Systems (DAPSYS 2002), Kluwer Academic Publishers, The Kluwer International Series in Engineering and Computer Science, Vol. 706, Linz, Austria, September 29-October 2, 2002
- [24] A. Iamnitchi and I. Foster, “On Fully Decentralized Resource Discovery in Grid Environments”, *Proceedings of an IEEE International workshop on Grid computing*, Denver, 2001
- [25] K. Czajkowski, S. Fitzgerald, I. Foster, and C. Kesselman, “Grid Information Services for Distributed Resource Sharing”, *Proceedings of the 10th IEEE International Symposium on High Performance Distributed Computing (HPDC-10)*, IEEE Press, 2001
- [26] I. Foster, et al., “A Distributed Resource Management Architecture that Supports Advanced Reservations and Co-Allocation”, *Proceedings of the International Workshop on Quality of Service*, 1999
- [27] K. Czajkowski, et al. “A resource management architecture for metacomputing systems”, *Proceedings of the 4th Workshop on Job Scheduling Strategies for Parallel Processing*, Springer-Verlag, LCNS 1459, 1998, pp. 62-82
- [29] P. Stelling, I. Foster, C. Kesselman, C. Lee, and G. von Laszewski, “A Fault Detection Service for Wide Area Distributed Computations”, *Cluster Computing*, Vol. 2, No. 2, 1999, pp. 117-128

Peer-to-Peer Computing

- [30] I. Foster and A. Iammitchi, “On Death, Taxes, and the Convergence of Peer-to-Peer and Grid Computing”, *Proceedings of the 2nd International Workshop on Peer-to-Peer Systems (IPTPS '03)*, February 20-21, 2003
- [31] M. Ripenau, “Mapping the Gnutella Network: Properties of Large-Scale Peer-to-Peer Systems and Implications for System Design”, TR-2001-26, University of Chicago, July 2001
- [32] A. Oram (ed.), Peer-to-Peer: Harnessing the Power of Disruptive Technologies, O'Reilly and Associates, Sebastopol, CA, 2001
- [33] M. Castro, et al., “One Ring to Rule them All: Service Discovery and Binding in Structured Peer-to-Peer Overlay Networks”, *The Proceedings of the Tenth ACM SIGOPS European Workshop*, ACM, September 22-25, 2002, Saint-Émilion, France
- [34] P. Keyani, B. Larson, and M. Senhil, “Peer Pressure: Distributed Recovery from Attacks in Peer-to-Peer Systems”, in *Web Engineering and Peer-to-Peer Computing*, Gregori, E. et al. (eds.), NETWORKING 2002 Workshops, Pisa, Italy, May 19-24, 2002, Revised Papers, Lecture Notes in Computer Science 2376 Springer 2002, ISBN 3-540-44177-8, pp. 306-320
- [35] D. C. Verma, et al. “SRIRAM: A scalable resilient autonomic mesh”, *IBM SYSTEMS JOURNAL*, VOL 42, NO 1, 2003, pp. 19-28
- [36] H-C Hsiao and C-T King, “Neuron – A Wide-Area Service Discovery Infrastructure”, *Proceedings of the International Conference on Parallel Processing (ICPP '02)*, August 18-21, 2002, p. 455
- [37] E. Halepovic and R. Deters, “JXTA Performance Model”, draft submitted paper from the Department of Computer Science at the University of Saskatchewan
- [38] S. Blanchfield, An Anonymous and Scaleable Distributed Peer-to-Peer System, B.A. Thesis, the University of Dublin, Trinity College, Ireland, 2002
- [39] J. Mischke and B. Stiller, Design Space for Distributed Search (DS)² – A System Designer's Guide, TIK-Report No. 151, Computer Engineering and Networks Laboratory, Swiss Federal Institute of Technology, Zurich, September 2002

Peer-to-Peer Computing (cont.)

- [40] S. Joseph, “NeuroGrid: Semantically Routing Queries in Peer-to-Peer Networks”, *Proceedings of the 1st International Workshop on Peer-to-Peer Systems (IPTPS'02)*, March 7-8, 2002
- [41] J. Lan, Cache Consistency Techniques for Peer-to-Peer File Sharing Networks, Master’s Thesis, Department of Computer Science, University of Massachusetts Amherst, June 26, 2002.
- [42] N. Blundell and L. Mathy, “An Overview of Gnutella Optimisation Techniques”, *Proceedings of Post Graduate Networking Conference 2002 (PGNet 2002)*, June 17-18, 2002
- [43] M. Barbosa, et al., “Using Locality of Reference to Improve Performance of Peer-to-Peer Applications”, draft paper submitted to the *4th International Workshop on Software Performance (WOSP 2004)*

Grid Simulators

- [44] R. Buyya and Manzur Murshed, “GridSim: a toolkit for the modeling and simulation of distributed resource management and scheduling for Grid computing”, *Concurrency and Computation: Practice and Experience*, Vol. 14, 2002, pp. 1175-1220
- [45] W. Bell, et al., “Simulation of Dynamic Grid Replication Strategies in OptorSim”, *Proceedings of the 3rd Int'l IEEE Workshop on Grid Computing (Grid'2002)*, Baltimore, 2002
- [46] I. Legrand and H. Newman, “The Monarc Toolset For Simulating Large Network-Distributed Processing Systems”, *Proceedings of the 2000 Winter Simulation Conference*, pp. 1794-1801, 2002
- [47] H. Casanova and L. Marchal, “A Network Model for Simulation of Grid Application”, INRIA Research Report No. 4596, October 2002
- [48] Aurora (a FreeNet Simulator) <http://www.doc.ic.ac.uk/~twh1/academic/>